International Finance: Putting Theory Into Practice

Piet Sercu
Leuven School of Business and Economics
Katholieke Universiteit Leuven

14:20 on 2 July 2008

Preface

About this book

This book had a forerunner—"International Financial Markets and The Firm", co-authored with Raman Uppal, which came out in 1995. By 2003 or 2004 Raman and I had agreed that a text full of Italian Lira or German Marks and where traders still had a full two minutes to respond to market makers' quotes, might sooner or later risk getting outdated. Starting the revision itself turned out to be much more difficult than agreeing on the principle, though. In the end Raman, being so much busier and more rational than I am, preferred to bow out. How right he was. Still, now that the effort has become a sunk cost, forever bygone, I find that episodes where I sincerely curse the book (and myself and Princeton University Press) are becoming fewer and farther between. Actually, there now are several passages I actually begin to like.

Like the previous book, the book still targets finance students, or at least students that want a genuine finance text, not an international-management or -strategy text with a finance slant nor an international monetary economics text with some corporate applications. There is a continued bias in favor of financial markets and economic logic; the aim is to provide students with a coherent picture of international markets and selected topics in multinational corporate finance. Sure, during everyday practice later on, this framework will then get amended and corrected and qualified; but the feeling of fundamental coherence will remain, we hope.

This book is more analytical than the modal text in the field. Compared to the Sercu-Uppal book, some of the math has been dropped and new matter has been added. As before, a lot of it is in Appendices, thus stressing its optional character. The main difference, I think, is that the in-text math is brought in differently. While in International Financial Markets we had every theorem or proof followed by an example, now the example comes first whenever that is possible. If so, the proof is often even omitted, or turned into a DoItYourself assignment. In fact, a third innovation is that, at least in the chapters or sections that are sufficiently analytical rather than just factual, the reader is invited to prove or verify claims and solve analogous problems. The required level of math is surely not prohibitive; anybody who has finished a good finance course should be able to master these DoItYourself assignments. Still, while the required level of mathematical prowess is low, a capacity for abstract thinking and handling symbols remains vital.

Every Part, except the Intro one, now has its own introductory case, which is intended to stimulate the reader's appetite and which can be a source of assignments. The cases usually cover issues from most chapters in the Part.

A fifth change is that the Part on exchange-rate pricing is much reduced. The former three Chapters on exchange-rate theories, predictability, and forward bias are now shrunk to two. And, lastly, three wholly new chapters have been added: two on international stock markets—especially crosslisting with the associated corporate-governance issues—and one on Value at Risk.

Typically, a preface like this one continues with a discussion and motivation of the book's content. But my feeling is that most readers—and surely students—skip prefaces anyway. Since the motivation of the structure is quite relevant, that material is now merged into the general introduction chapter, Chapter 1.

How to use this book

The text contains material for about two courses. One possibility is to take the second Part, *International Financial Markets*, as one course, and group the more business-finance oriented material (grouped into *Exchange Risk, Exposure, and Risk Management* (III) and *Long-Term Financing and Investments* (IV)) as a second. Fixed-income markets, which now is in Part III, could be included in the markets/instruments course, like it was in the 1995 book; and the whole package can also duplicate as an intro derivatives course, along with the apocryphal Chapter **??** that is available on my website. I myself run two 40-hr courses covering, respectively Parts II-III (Instruments, Risk Management) and Part IV (Stocks, bonds, capital budgeting).

For one single course one could focus, in Part II, on spot (Chapter 3) and forwards (Chapters 4 and 5), and then continue with the chapters on relevance of hedging and exposure (Chapters 12 and 13), to finish with capital budgeting (Chapter 21); this shortlist can be complemented by a few chapters of your liking.

Leuven, December 2006.

About the author

Piet Sercu is Professor of International Finance at the *Katholieke Universiteit Leuven.* He holds the degrees of Business Engineer, Master of Business Administration, and Doctor in Applied Economics from K.U. Leuven. He taught at the Flemish Business School in Brussels (1980-1986), prior to returning to Leuven, where he currently teaches the International Business Finance courses in the Masters and Advanced Masters programs. He also held Visiting Professor appointments at New York University, Cornell University, the University of British Columbia, the London Business School, and *Université Libre de Bruxelles.* He taught shorter finance courses in Helsinki, Bandung (Indonesia), Leningrad, and India (as an UNDP expert and, in 1994, as a fellow of the European Indian Cooperation and Exchange Programme), and regularly teaches executive courses. He held the 1996/7 Francqui Chair at the *Facultés Universitaires Notre-Dame de la Paix* at Namur, and the 2000/04 PricewaterhouseCoopers Chair on Value and Risk at KU Leuven, together with Marleen Willekens. Until 2000, he organized and taught doctoral courses in the European Doctoral Education Network, as part of the Finance faculty of the European Institute for Advanced Studies in Management. He was the 1994 Vice-President and 1995 President of the European Finance Association, won the 1999 Western Finance Association award for Corporate Finance (with Xueping Wu and Charley Park) and was Hanken Fellow in 2002.

His early research focused on International Asset Pricing with real exchange risk and inflation risk. He also did some work on corporate take-over models and lending but has recently returned to International Finance and hedging. He has published in the Journal of Finance, Journal of Banking and Finance, Journal of International Money and Finance, European Economic Review, and other journals. He is on the editorial boards of the European Financial Management Journal and the Journal for International Financial Markets, Institutions and Money.

Piet Sercu and Raman Uppal jointly won the 1995 Sanwa Prize for a monograph in International Finance, *Exchange Rate Volatility, Trade, and Capital Flows under Alternative Currency Regimes*, published by Cambridge University Press in 2000 and 2006. They also have produced *International Financial Markets and The Firm* (International Thomson Publishers, Cincinnati-London, 1995), the forerunner to this book and the source of much of its material. There are also a number of joint academic articles.

Acknowledgments

There are many individuals who played an important role in the production of this book. First and foremost I thank Raman Uppal, not only for his invaluable contribution to the first book but also for the discussions about how to structure a new version, for translating the text parts of the old manuscript into LaTeX, and setting up a master file system to produce the whole. Thanks also to the former and current doctoral students or assistants who read earlier drafts of the first and second book and suggested several improvements: Badrinath H. R., Thi Ngoc Tuan Bui, Katelijne Carbonez, Cédric de Ville de Goyet, Kathy Dehoperé, Marian Kane, Fang Liu, Rosanne Vanpée, Tom Vinaimont, and Xueping Wu. Marian, especially, did lots of work on the exercises, and Katelijne on the revised text. Prof. Martina Vandebroek occasionally helped with spreadsheets and graphs, at which she Excels; she, Fang Liu and Badrinath H. R. provided the empirical results for Chapters 10 and 11. Many thanks, lastly, to colleagues who read drafts and provided comments, some of them at Princeton's request but some even of their free will: Hu Shengmei, Karen Lewis, Bernard Dumas, Stan Standaert, Charles van Wymeersch, and an anonymous (but wholly positive) referee reporting to Princeton UP. Of course, I remain responsible for all remaining errors. Comments and feedback from readers about errors, presentation, and contents are very welcome: do email to `piet.sercu@econ.kuleuven.be`.

I dedicate this book to my parents, Jan Sercu and the late Térèse Reynaert, and to my wife Rita and children Maarten and Jorinde, who have patiently put up with my inattentive absent-mindedness during the time it has taken to complete this project and, come to think of it, most of the time before and after.

December 2007

Contents

Part I

Introduction and Motivation for International Finance

Chapter 1

Why does the Existence of Borders Matter for Finance?

Almost tautologically, international finance selects from the broad field of finance those issues that have to do with the existence of many distinct countries. The fact that the world is organized into more or less independent entities instead of a single global state complicates a CFO's life in many ways—ways that matter far more than does the existence of provinces or states or *Landen* or *départements* within a country. Below, we discuss

- the existence of national currencies and, hence, the issue of exchange rates and exchange risk;
- the segmentation of goods markets along predominantly national lines; in combination with price stickiness, this makes most exchange-rate changes "real";
- the existence of separate judicial systems, which further complicates the already big issue of credit risk, and has given rise to private-justice solutions;
- the sovereign autonomy of countries, which adds political risks to standard commercial credit risks
- the existence of separate and occasionally incompatible tax systems, giving rise to issues of double and triple taxation.

We review these items in Section 1. Other issues or sources of problems, like differences in legal systems, investor protection, corporate governance, and accounting systems are not discussed in much depth—not because they would be irrelevant but for the simple reasons that there is too much heterogeneity across countries and I have no expertise in them. Still, there is a chapter that should create a basic awareness in these issues, so that the reader can then critically look at the local regulation and see its relative strengths and weaknesses,

The above list includes some of the extra issues a CFO in an international company needs to handle when doing the standard tasks of funding, evaluation, and risk management (Section 2). The outline of how we will work our way through all this matter follows in Section 3.

1.1 Key Issues in International Business Finance

1.1.1 Exchange-rate Risk

Why do most countries have their own money? One disarmingly simple reason is that printing bank notes is profitable, obviously, and even the minting of coins is usually a positive-NPV business. In the West, at least since the days of the Greeks and Romans, governments have been involved as monopoly producers of coins or at least as receivers of a royalty ("seignorage") from the use of the official logo. More recently, the ascent of paper money, where profit margins are almost too good to be true, has led to official monopolies virtually everywhere. One reason why money production is not handed over to the UN or the IMF or WB is that governments dislike giving up their monopoly rents. For instance, the shareholders of the European Central Bank are the individual Euro-countries, not the EU itself; that is, the countries have given up their monetary independence, but *not* their seignorage. In addition, having one's own money is a matter of national pride too: most Brits or Danes would not even dream of surrendering their beloved Pound Sterling or Crown for, of all things, a European currency. Lastly, a country with its own money can adopt a monetary policy of its own, tailored to the local situation. Giving up a local policy was a big issue at the time the introduction of a common European money was being debated.[1]

If money had intrinsic value (e.g. a silver content), if that intrinsic value were stable and immediately obvious to anybody, and if coins could be de-minted into silver and silver re-minted into coins at no cost and without any delay, then the value of a German Joachimsthaler relative to a Dutch Florin and a Spanish Real would all be based on their relative silver content, and would be stable. But in practice, many sovereigns were cheating with the silver content of their currency, and got away with it in the short run. Also, there are costs in identifying a coin's true intrinsic value and in converting Indian coins, say, into Moroccan ones. Finds of hoards dating from

[1]Following a national monetary policy assumes that prices for goods & services are sticky, that is, do not adjust quickly when money supply or the exchange rate are being changed. (If prices would fully and immediately react, monetary policy would not have any 'real' effects). Small open economies do face the problem that local prices adjust too fast to the level of the countries that surround them. So it's not a coincidence that Monaco, San Marino, Andorra and the Vatican don't bother to create their own currencies. Not-so-tiny Luxembourg similarly formed a monetary union with Belgium as of 1922. Those two then fixed their rate to the DEM and NLG with a 1% band in 1982. For more countries that gave up, or never had, an own money see Wikipedia, Monetary Union. See the section on Currency Boards, in this chapter, about countries that give up monetary policy but not seignorage.

Figure 1.1: **Relative prices of the Big Mac across the World, May 2006**

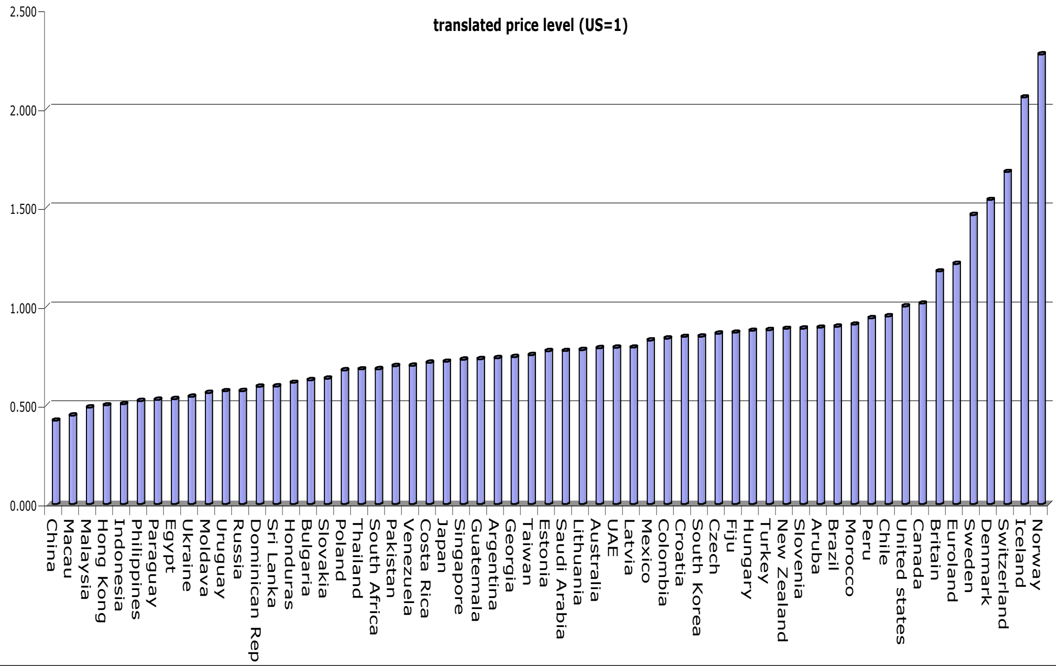

Source Based on data from *The Economist*, May 26, 2006

the Roman or Medieval times reveal astounding differences in the silver content of various coins with the same denomination. For instance, among *solidus* pieces from various mints and of many vintages, some have silver contents that are twice that of other solidus coins found in the same hoard. In short, intrinsic value did never nail down the market value in a precise way, not even in the days when coins really were made of silver, and as a result exchange rates have always fluctuated. Since the advent of paper money and electronic money, of course, intrinsic value no longer exists: the idea that paper money was convertible into gold coins lost all credibility after WW1. After WW2, governments for some time controlled the exchange rates, but largely threw in the towel in 1973-4. Since then, exchange rates are based on relative trust, a fickle good, and the resulting exchange-rate risk is a fact of life for all major currency pairs.

Exchange risk means that there is uncertainty about the value of an asset or liability that expires at some future point in time and is denominated in a foreign currency ("*contractual exposure*"). But exchange risk affects a company's financial health also via another channel—an interaction, in fact, with another inter-national issue: segmentation of the consumption goods markets.

1.1.2 Segmentation of the Consumer-good Markets

While there are true world markets—and, therefore, world prices—for commodities, many consumer goods are really priced locally, and for traditional services the international influence is virtually absent. Unlike corporate buyers of say oil or corn or aluminum, private consumers do not bother to shop around internationally for the best prices: the amounts at stake are too small, and the transportation cost and hassle and delay from international trade would be prohibitive anyway. Distributors, who are better placed for international shopping-around, prefer to pocket the resulting quasi-rents themselves rather than passing them on to consumers. For traditional services, international trade is not even an option. So prices are not homogenized internationally even after conversion into a common currency. One strong empirical regularity is that, internationally, prices rise with GDP/capita. In Figure 1.1, for instance, you see prices of the Big Mac in various countries, relative to the US price. Obviously, developed countries lead this list, with growth countries showing up as less expensive by *The Economist*'s Big Mac standard. The ratio of Big Mac prices Switzerland/China is 3.80. Norway (not shown here) was even more than five times more expensive than China, in early 2006; and two years before, the gap Iceland/South Africa was equally wide.

Within a country, by contrast, there is less of this price heterogeneity. For example, price differences between "twin" towns that face each other across the US-Canadian or US-Mexican border are many times larger than differences between East- and West-coast towns within the US. One likely reason that contributes to more homogenous pricing within a country is that distributors are typically organized nationally. Of course, the absence of hassle with customs and international shippers and foreign indirect tax administrations also helps.

A second observation is that prices tend to be sticky. Companies prefer to avoid price increases, because the harm done to sales is not easily reversed: consumers are resentful, or they just write off the company as "too expensive" so that they do not even notice when prices come down again. Price decreases, on the other hand, risk setting off price wars, and so on.

Now look at the combined picture of (*i*) price stickiness, (*ii*) lack of international price arbitrage in consumption-good markets, and (*iii*) exchange-rate fluctuations. The result is *real exchange risk.* Barring cases of hyperinflation, short-run exchange-rate fluctuations have little or nothing to do with the internal prices in the countries that are involved. So the appreciation of a currency is not systematically accompanied by falling prices abroad or soaring prices at home so as to keep goods prices similar in both countries. As a result, appreciation or depreciation can make a country less attractive as a place to produce and export from or as a market to export to. They therefore affect the market values and competitiveness of companies and economies ("*economic exposure*"). For instance, the soaring USD in the Reagan years has meant the end of many a US company's export business, and the rise of the DEM in the 70s forced Volkswagen to become a multi-country producer.

Real exchange risk also affects asset values in a more subtle way. Depending on where they live, investors from different countries realize different real returns from one given asset if the real exchange rate changes. Thus, one of the fundamental assumptions of *e.g.* the CAPM, that investors all agree on the returns and risks of all assets, becomes untenable. While this may sound like a very theoretical issue, it becomes more important once you start thinking about capital budgeting. For instance, a US firm may be considering an investment in South Africa, starting from projected cash flows in South-African Rand (SAR). How to proceed? Should the managers discount them using a SAR discount rate, the way a local investor would presumably do it, and then convert the PV into USD using the current spot rate? Or should they do it the US way: use expected future spot rates to convert the data into expected USD cash flows, to be discounted at a USD rate? Should both approaches lead to the same answer? *Can* they, in fact?

Exchange risk is the issue that takes up more space than any other separate topic in this book. Its importance can be seen from the fact that so many instruments exist that help us cope with this type of uncertainty: forward contracts, currency futures and options, and swaps. You need to understand all these instruments, their interconnections, their uses and limitations, and their risks.

1.1.3 Credit risk

If a domestic customer does not pay, you resort to legal redress, and the courts enforce the ruling. Internationally, one problem is that at least two legal systems are involved, and they may contradict each other. Usually, therefore, the contract will stipulate what court will rule and on the basis of what law—say Scottish law in a New York court (I did not make this up). Even then, the new issue is that this court cannot enforce its ruling outside its own jurisdiction.

This has given rise to private-contract solutions: we seek guarantees from specialized financial institutions (banks, factors, insurance companies) that (*i*) are better placed to deal with the credit risks we shifted towards them, and (*ii*) have an incentive to honor their own undertakings because they need to preserve a reputation and safeguard relations with fellow banks etc. So you need to understand where these perhaps Byzantine-sounding payment options (like D/A, D/P, L/C without or with confirmation, factoring, and so on) come from, and why and where they make sense.

1.1.4 Political risk

Governments that decide or rule as sovereigns, having in mind the interest of their country (or claiming to have this in mind), cannot be sued in court as long as what they do is constitutional. Still, these decisions can hurt a company. One example is imposing currency controls, that is, block some or all exchange contracts, so that the money you have in a foreign bank account gets stuck there (*transfer risk*). You need to know how you can react pro- and retroactively. You also need to know how this risk has to be taken into account in international capital budgeting. If and when your foreign-earned cash flow gets stuck abroad, it is obviously worth less than its nominal converted value because you cannot spend the money freely where and how you want—but how does one estimate the probabilities of this happening at various dates, and how does one predict the size of the value loss?

Another political risk is expropriation or nationalization, overtly or on the stealth. While governments can also expropriate locally-owned companies (like banks, in 1981 France), foreign companies in the "strategic" sectors (energy, transportation, mining & extraction, and, flatteringly, finance) are especially vulnerable: most of them were expropriated or had to sell to locals in the 1970s. The 2006 Bolivian example, where President Evo Morales announced that "The state recovers title, possession and total and absolute control over [our oil and gas] resources" (*The Economist*, May 4, 2006.) also has to do with such a sector. Again, one issue for the finance staff is how to factor this in into NPV calculations.

1.1.5 Capital-Market Segmentation Issues, including Aspects of Corporate Governance

A truly international stock and bond market does not exist. First, while stocks and bonds of big corporations do get traded in many places and are held by investors all over the world, mid-size or small-cap companies are largely one-country instruments. Second, portfolios of individual and institutional investors exhibit strong home bias—that is, heavy overweighting of local stocks relative to foreign stocks—even regarding their holdings of shares in large corporations. A third aspect of fragmentation in stock markets is that we see no genuine international stock exchanges (in the sense of institutions where organized trading of shares takes place); instead, we have a lot of local *bourse*s. A company that wants its shares to be held in many places gets a listing on two or three or more exchanges (*dual* or *multiple listings*; *cross-listing*): being traded in relatively international places like London or New York is not enough, apparently, to generate worldwide shareholdership. How come?

The three phenomena might be related, and caused by the problem of asymmetric information and investor protection. Valuing a stock is more difficult than valuing a bond, even a corporate bond, and the scope for misrepresentation is huge, as the railroad and dotcom bubbles have shown. All countries have set up some legislation and regulation to reduce the risks for investors, but there are enormous differences in the amount of information, certification and vetting required for an initial public offering (IPO). All countries think, or claim to think, the other countries are fools by imposing so much/little regulation. The scope for establishing a common world standard in the foreseeable future is nil. Pending this, there can be no single world market for stocks.

The same holds for disclosure requirements once the stock has been launched, and the whole issue of corporate governance. The big issue here is how to avoid managers selfdealing or otherwise siphoning off cash that ought to be the shareholders'. Good governance systems contain checks and balances, like separation of the jobs of chairman of the Board of Directors and CEO; a sufficient presence of independent directors on the Board; an audit committee that closely watches the accounts; comprehensive information provision towards investors; a willingness, among the board members, to fire poorly performing CEO's, perhaps on the basis of pre-set performance criteria; a board that can be fired by the Assembly General Meeting in one shot (as opposed to staggered boards, where every year only one fifth comes up for (re)election, for example); and a AGM that can formulate binding instructions to the Board and the CEO. Good governance also requires good information provision, with detailed financial statements accompanied by all kinds of qualitative information.

But governance is not just a matter of corporate policies: it can, and ideally must, be complemented by adequately functioning institutions in the country. For instance, how active and independent are auditors, analysts (and, occasionally, newspaper reporters)? Is a periodic evaluation of the company's financial health by its house bank(s), each time loans are rolled over or extended, a good substitute for outside scrutiny? Are minority shareholders well protected, legally? How stringent are the disclosure and certification requirements, and are they enforced? Are there active large shareholders, like pension funds, that follow the company's performance and put pressure onto management teams they are unhappy with? Is there an active market for corporate officers, so that good managers get rewarded and (especially) *vice versa*? Is there an active acquisition market where poorly performing companies get taken over and reorganized? Again, on all these counts there are huge differences across countries, which makes it impossible to set up one world stock market. The OECD has been unable to come up with a common stance on even something as fundamental as accounting standards. Telenet, a company discussed in a case study in Part IV, has three sets of accounts: Belgian GAAP, US GAAP, and IFRS. Even though in the US its shares are only sold to large private investors rather than the general public, Telenet still had to create a special type of security for the US markets.

In short, markets are differentiated by regulation and legal environment. In addition, companies occasionally issue two types of shares: those available for residents of their home country, and unrestricted stocks that can be held internationally. Some countries even impose this by law. China is a prominent example, but the list used to include Korea, Taiwan, and Finland/Sweden/Norway. Typically, only a small fraction of the shares was open to non-residents. Other legislation that occasionally still fragments markets is a prohibition to hold forex; restrictions or prohibitions on purchases of forex, especially for financial (i.e. investment) purposes; caps on the percentage of mutual funds or pension funds invested abroad, or minima for domestic investments; dual exchange rates that penalize financial transactions relative to commercial ones; taxes on deposits by non-residents; requirements to invest at zero interest rates at home, proportionally with foreign investments or even with imports, and so on—you name it.

In OECD countries or NICs, this type of restrictions is now mostly gone. In December 2006, Thailand imposed some new regulations in order to discourage inflows—usually the objective is to stop outflows—but hastily reversed them after the Bangkok stock market had crashed by 15 percent; this example goes to show that this type of restriction is simply *not done* anymore. But some countries never lifted them altogether, like Chile, while in other countries the bureaucratic hassle is still strongly discouraging (India) or virtually prohibiting (Russia) capital exports.

There are two repercussions for corporate finance. One is via the shareholders. Specifically, in countries with serious restrictions on outward investments, the investment menu is restricted and different from the opportunity set available to luckier investors elsewhere. This then has implications for the way one works with the CAPM: companies in a walled-off country have to define the market portfolio in a strictly local way, while others may want to go all the way to the world-market version of the market portfolio. So companies' discount rates are affected and, therefore, their direct investment decisions. Another corporate-finance implication is that a company that wants to issue shares abroad cannot simply go to some "international" market: rather, it has to select a country and, often, a segment (an exchange—which exchange? which board?—or the over-the-counter market or the private-investors segment), carefully weighting the costs and benefits of its choices. An important part of the costs and benefits have to do with the corporate-governance and disclosure ramifications of the country and market segment one chooses.

1.1.6 International Tax Issues

Fiscal authorities are understandably creative when thinking up excuses to tax. For instance, they typically want to touch all residents for a share in their income, whether that income is domestic or foreign in origin; but they typically also insist on taxing anybody making money inside the territory, whether the earner is a resident or not. So a Icelandic professor making money in Luxembourg as visiting faculty would be taxed by both Luxembourg—she did make money there—and by Iceland—she is a resident there.

In corporate examples things get even worse. When an Icelandic corporation sets up shop in Luxembourg, the subsidiary is taxed there on its profits: it is a resident of Luxembourg, after all. But when that company then pays a dividend to its parent, both Luxembourg and Iceland may want to tax the parent company again: the parent makes money in L, but is a resident in I.

Fortunately, legislators everywhere agree that double or triple taxation maybe somewhat overdoing things, so they advocate neutrality. But, as we shall see, there is no agreement as to how a "neutral" system can be defined, let alone how it is to be implemented. This makes life for the CFO complicated. But it also makes life exciting, because of the loopholes and clever combinations ("treaty shopping") that can substantially affect the tax burden.

1.2 What is on the International CFO's desk?

This book is a text on international finance. Thus, it does not address issues of multinational corporate strategy, and the discussion of international macroeconomics is kept to a minimum. Within the finance discipline, it addresses only the problems caused by the existence of many countries, as described in the preceding section.

One way to further describe the material is to think about the tasks assigned to an international financial manager. These tasks include asset valuation, international funding, the hedging of exchange risk, and management of other risks. We hasten to add that these functions cannot be viewed in isolation, as will become clear as we proceed.

1.2.1 Valuation

One task of an international finance officer is the valuation of projects with cash flows that are risk free in terms of the foreign currency. For example, the manager may need to evaluate a large export order with a price fixed in foreign currency and payable at a (known) future date. The future cash flow is risky in terms of the domestic currency because the future exchange rate is uncertain. Just like one would do with a domestic project with cash flows that are risky in terms of the domestic currency, this export project should be subject to a Net Present Value (NPV) analysis. Thus, the manager needs to know how to compute present values when the source of risk is the uncertainty about the future exchange rate. Valuation becomes even more complicated in the case of foreign direct investment (FDI), where the cash flows are random even in terms of the foreign currency. The issues to be dealt with now are how to discount cash flows subject to both business risk and exchange risk, how to deal with tax complications and political risks inherent in FDI, and how to determine the cost of capital depending on whether or not the home and foreign capital markets are segmented.

1.2.2 Funding

A second task is, of course, funding the project. A standard financing problem is whether the firm should issue equity, debt, or equity-linked debt (like convertible bonds). If bonds are issued or a loan is taken out, the standard questions are what the optimal maturity is, and whether the terms offered by a bank or a group of banks are attractive or not. In an international setting, the additional issue to be considered is whether the bond or loan should be denominated in home currency or in another one, whether or when there are any tax issues in this choice, how the risk can be quantified when it is correlated with other risks, and so on.

If funding is done in the stock markets, the issue is whether to issue stocks locally or to get a secondary listing elsewhere—or perhaps even move the company's primary listing abroad. The targeted foreign market may be better organized, have more analysts that know and understand your business, and give access to deep-pocketed investors who, being well-diversified already, are happy with lower expected returns than the current shareholders. But there are important corporate-governance issues as well, as we saw: getting a listing in a tough place is like receiving a certificate of good behavior and making a strong commitment to behave well in future too. So the mere fact of getting such a listing can lift the value of the company as a whole. There are, of course, costs too: publishing different accounts and reports to meet diverging accounting and disclosure rules can be cumbersome and expensive, and listing costs are not trivial either. Because of the corporate-governance issues, cross-listings are not purely technical decisions that belong to the CFO's competence: the whole Board of Directors should be involved.

1.2.3 Hedging and, more Generally, Risk Management

Another of the financial manager's tasks usually is to reduce risks, like exchange risk, that arise from corporate decisions. For example, a manager who has accepted a large order from a customer, with a price fixed in foreign currency and payable at some (known) future point in time, may need to find a way to hedge the resulting exposure to exchange rates.

There are, however, many other sources of uncertainty besides exchange rates. Some are also "market" risks: uncertainties stemming from interest rates, for instance, or commodity prices or, for some companies, stock market gyrations. Exchange risk cannot be hedged in isolation, for the simple reason that market risks tend to be correlated. As a result, many companies want to track the remaining uncertainties of their entire portfolio of activities and contracts. This is usually summarized in a number called *value at risk* (VaR), the maximum loss that can be sustained with a given probability (say, 1 percent) over a given horizon (say, one day), taking into account the correlations between the market risks.

1.2.4 Interrelations Between Risk Management, Funding and Valuation

While the above taxonomy of CFO assignments is logical, it does not offer a good structure for a textbook. One reason is that valuation, hedging, and funding are interrelated. For instance, a firm may be unwilling to accept a positive-NPV export contract (valuation) unless the currency risk can be hedged. Also, the funding issue cannot be viewed in isolation from the hedging issue. For example, a Finnish corporation that considers borrowing in Yen, should not make that decision without pondering how this loan would affect the firm's total risk. That is, the decision to borrow Yen may be unacceptable unless a suitable hedge is available. In another example, a German firm that has large and steady dollar revenues from exports might prefer to borrow USD because such a loan provides not just funding, but also risk reduction. In short, project evaluation, funding, and hedging have to be considered together.

But risks do not stop at market risks. There are credit risks, political risks, operational risks, reputation risks, and so on, and also these interact with the more financial issues. For instance, the evaluation of an export project should obviously take into account the default risk. Similarly, NPV computations for FDI projects should account for the risk that foreign cash flows may be blocked or that the foreign business may be expropriated.

1.3 Overview of this Book

In the preceding section, we discussed the key issues in international finance on the basis of managerial functions. As said, this is not a convenient way to arrange the text because the functions are all interlinked. Instead, we proceed as follows. We begin with an introductory chapter on the history of the international monetary system. The remainder of this textbook, then, is divided into four parts: (*II*) International financial markets and instruments; (*III*) Exchange rate risk, exposure, and risk management; and (*IV*) Long-term financing and investment decisions. In most of the chapters except the next one, the focus is on corporate financial issues, such as risk management and funding and capital budgeting. Let us briefly review the contents of each part below.

1.3.1 Part I: Motivation and Background Matter

After the present motivational chapter, we go over some background material: how is money created, how is it paid internationally, what is the role of governments in exchange markets, and what does the Balance of Payments mean for a country?

1.3.2 Part II: International Financial Markets

Part II of the book describes the currency market in its widest sense, that is, including all its satellites or derivatives. Chapter 3 describes spot markets. Forward markets, where price and quantity are contracted now but delivery and payment take place at a known future moment, are introduced in Chapter 4, in a perfect-markets setting. Chapter 5 shows how and when to use contracts in reality: for arbitrage, taking into account costs; for hedging; for speculation; and for shopping-around and structured-finance applications including, especially, swaps. Currency futures and modern currency swaps, both of which are closely related to forward transactions, are discussed in Chapters 6 and 7, respectively. Chapter 8 introduces currency options and shows how these options can be used to hedge against (or alternatively, speculate on) foreign exchange risk. How one can price currency options is explained in Chapter 9; we mostly use the so-called binomial approach but also link it to the famous Black-Merton-Scholes model.

At any instant, the market value of a forward, futures, or options contract depends on the prevailing spot rate (and, if the contract is not yet at the end of its life, also on the domestic and foreign interest rates). This dependence on the future spot rate means that all these contracts can be used to hedge the exchange-rate risk to which the firm is exposed. The dependence of these contracts on the future spot rate also means that their current market values can be expressed, by relatively simple arbitrage arguments, as functions of the current spot rate and of the domestic and foreign interest rate. Throughout this part of the text, a unified approach based on arbitrage-free pricing is used to value these assets whose payoffs are dependent on the exchange rate.

1.3.3 Part III: Exchange Risk, Exposure, and Risk Management

This parts opens with a discussion of the behavior and predictability of nominal and real exchange rates (Chapters 10 and 11). We conclude that exchange rates are hard to explain, let alone to predict, and that most of the nominal uncertainty is also real, thus affecting the long-term value of a company.

This may sound like a good excuse to hedge. Yet one could argue that (*i*) hedging is a standard financial transaction; (*ii*) in efficient markets, financial transactions are zero-NPV deals; (*iii*) therefore, hedging does not add value. In Chapter 12 we show the way out of this fallacy: hedging does add value if it does more than just increase or decrease the firm's bank account—that is, if and when it affects the firm's operations. Given that firms may want to hedge, the next issue is how much to hedge: what is the size of the exposure (Chapter 13)? We distinguish between contractual, operational, and accounting exposures. Value at risk is reviewed in Chapter 14. Chapter 15 concludes this part with a description and critical discussion of the various ways to insure credit risks and transfer risks in international trade.

1.3.4 Part IV: Long-term Financing and Investment Decisions

The prime sources for long-term financing are the markets for fixed-interest instruments (bank loans, bonds) and stocks. We review the international aspects of these in Chapters 16 and 17-18, respectively, including the fascinating issue of cross-listing and corporate governance. Expected returns on stocks provide one key input of investment analysis, so in Chapter 19 we consider the CAPM and the adjustments to be made to take into account real exchange risk. The other inputs into NPV computations are expected cashflows, and these are typically quite similar to what one would see in domestic projects. There is one special issue here, international taxes (Chapter 20). In Chapter 21 we see how to do the actual NPV, extending the usual two-step approach—NPV followed by Adjusted NPV to take into account the aspects of financing, relevant in imperfect markets—to a three-step version to separately handle intra- and extra-company financial arrangements. We conclude with an analysis of joint-venture projects, where NPV is mixed with the issue of designing a fair profit-sharing contract (Chapter 22).

Here we go, then.

Chapter 3

Spot Markets for Foreign Currency

In this chapter, we study the mechanics of the spot exchange market. The first section explains the various ways in which exchange rates can be quoted, and the second section how the exchange markets themselves operate. Section 3 then considers exchange transactions in greater detail, focusing on bid and ask rates (that is, the rates at which a bank buys and sells). This also gives us an opportunity to learn about arbitrage. Specifically, in the third section, we shall already apply arbitrage arguments to the simplest possible problem, the relation between rates quoted by different banks for the same currency. Understanding this simple application now will make it easier to digest more complicated versions of similar arguments later. One such application already occurs in the fourth section, where we use arbitrage arguments to explain how exchange rates quoted, for example, by German banks (against EUR) relate to rates offered by New Zealand banks (against the NZD).

The chapter ends with the concepts of, and empirical evidence on, "Purchasing Power Parity (PPP)" rates and real exchange rates. The conclusion of that part will be that exchange rates can make or break an exporting company, not just because of capital losses on foreign-currency-denominated receivables but possibly also because of a loss of competitiveness. Exchange risk even interferes with capital market equilibrium and the capital asset pricing model. These findings motivate the attention given to exchange rates in this book.

3.1 Exchange Rates

As we begin exploring exchange rates, we first provide a definition. We then describe the convention that is used to quote exchange rates throughout this book, as well as the conventions used in the exchange market. Finally, we explain how exchange rates are quoted in the presence of bid-ask spreads.

3.1.1 Definition of Exchange Rates

An exchange rate is the amount of a currency that one needs in order to buy one unit of another currency, or it is the amount of a currency that one receives when selling one unit of another currency. An example of an exchange rate quote is 0.8 USD per CAD (which we will usually denote as "USD/CAD 0.8"): you can, for instance, buy a CAD by paying USD 0.80.

In the above, we have combined currency names following the conventions in physics: EUR/USD means euros per dollar just like km/hr means kilometers per hour. This is the most logical. For instance, if you exchange 3m dollars into euros at a rate of 0.8 euros per dollar, the result is a number of euros. This fits with our notation:

$$\text{USD } 3m \times \text{EUR/USD } 0.8 = \text{EUR } 2.4m. \tag{3.1}$$

This may seem self-evident. The reason why we bring this up is that pros do it differently. In the convention typically adopted by traders, bankers and journalists, EUR/USD is not the dimension of the quote but the name of the exchange rate: the Euro quoted in dollars, not its dimension. That is, traders etc. write "EUR/USD = 1.2345" whereas we write "$S_t =$ USD/EUR 1.2345". The dimension the trader asks for is USD/EUR, the inverse of what they write—but they do not mean a dimension, they mean a name.[1] In all our examples *we* use dimensions. The "name" notation pops up occasionally in press clippings or pictures of trading screens etc, and should not be a problem. To harden yourself, stare at the following line for a full minute:

$$\underbrace{\text{EUR/USD}}_{\text{name}} : \underbrace{\text{USD/EUR}}_{\text{dimension}} 1.25. \tag{3.2}$$

The tell-tale difference is that the dimension is immediately followed (or, occasionally, preceded) by the number. If there is no number, or if there is an "=" or "is" or "equals" etc between the ratio and the number, it must be the name of a rate. Sometimes practitioners drop the slash in the name and write EURUSD or EUR:USD instead of EUR/USD, which makes more sense.

It is even more crucial that you understand how exchange rates are quoted. While the notation is occasionally confusing—are we using dimensions or names?—there could be even more confusion as to which currency should be used as the numeraire. While you are familiar with the idea of buying goods and services, you may be less used to buying money with money. With exchange transactions, you need to agree which money is being bought or sold. There would be no ambiguity if one of the currencies were your home currency. A purchase then means that you obtained foreign currency and paid in home currency, the way you would do it with your other

[1] It is sometimes whispered that the trader notation comes from a kind of pseudo math like "EUR 1 = USD 1.2345", where one then "divides both sides by USD". The mind boggles. This is like denoting a speed as "1 hr = 100 km" instead of $v = 100$ km/hr.

purchases too; and a sale means that you delivered foreign currency and received home currency. If neither currency is your home currency, then you need to establish which of the two acts as the home currency.

Example 3.1

In a Paris bank, a tourist hands over USD 1,000 to the bank clerk and receives CAD 1,250 in return. This event would be described differently depending on whether the person is a US tourist, a Canadian, or a Frenchman:

- The US tourist would view this as a purchase of CAD 1,250 at a total cost of USD 1,000, implying a unit price of [USD 1,000]/[CAD 1.250] = USD/CAD 0.8.
- The Canadian would think of this transaction as a sale of USD 1,000 for CAD 1,250, implying a unit price of [CAD 1,250]/[USD 1,000] = CAD/USD 1.25.
- The Frenchman would regard this as an exchange of two foreign currencies, and would be at a loss if he would be asked which of these is being sold and which bought.

Among pros, the currency in which the price is expressed is called the *quoting currency*, and the currency whose price is being quoted is called the *base currency* or *reference currency*. We avoid the terms, except in the next two lines. We just noted that pros denote a rate as base/quoting (or, better, base:quoting) while its dimensions are quoting/base. A different issue is whether the quoting currency is the home or the foreign one.

3.1.2 Our Convention: Home Currency per Unit of Foreign Currency

Once we agree which country is, or acts as, the home country, we can agree to quote exchange rates as the price in units of home currency (HC), per unit of foreign currency (FC). That is, we quote the rate as HC/FC throughout this text, meaning that one unit of foreign currency is worth N home-currency units (dimension HC/FC). As we shall see, some people do it differently and state that with one unit of home currency, they can buy $M = 1/N$ units of foreign currency (FC/HC). We adopt the HC/FC convention because it is the most natural one. It is the convention we use when buying goods. For example, we say "the price is 5 dollars per umbrella" (HC/umbrella) not, "with one dollar you can buy one-fifth of an umbrella" (umbrellas per unit of home currency).

Example 3.2

1. A quote like USD/EUR 1.25 is an American's natural quote for the EUR; it is the USD price an American gets or pays per EUR. For Germans or other Eurolanders, a quote as EUR/USD (euros per dollar) is the more natural one.
2. A quote like USD/CAD 0.75 is an American's natural quote for the CAD, since the CAD is the currency in the denominator: a price in USD per CAD

Expressing prices in HC is the convention for not just umbrellas but also for financial assets. Thus, standard finance results hold: the current market value is the expected future value (including interest earned), discounted at a rate that takes into account the risk. Under the alternative quotation, confusingly, the current value would be determined by the inverse of the expected inverse of future value, multiplied by unity plus the required return. (If you just felt you had to read this sentence twice, you may want to consider reading end-of-chapter Teknote 3.1 instead.)

The direct (HC/FC) quoting convention used to be standard in continental Europe, and is called the direct quote, or the "right" quote. In the US, a price with dimension USD/FC is called "American terms". The alternative is called the "indirect" or "left" quote or, in the US, "European terms". Let's see who uses which and why.

3.1.3 The Indirect Quoting Convention

One group of people using mostly indirect quotes are professional traders in the US. Between 1944 and the mid-80s, each and every exchange deal went through the USD; even when a German needed to buy CHF, the DEM would first be converted into USD and these dollars were then exchanged for CHF. Naturally, when NY traders talk to, say, their German counterparts, both must talk the same language, quotewise; otherwise too much time would be wasted inverting each other's rates all the time. Both Germans and Americans actually preferred to quote in terms of DEM/USD rather than USD/DEM, for the simple reason that the official parities, set by the German government, were expressed in DEM/USD.[2] More in general, US professionals use the exchange-rate convention as quoted in the other country. Thus, for countries that quote directly themselves, like Japan, New York traders would talk JPY/USD. But in the case of countries that quote indirectly themselves, like the UK, pros would also use USD/GBP. Thus, US pros use indirect quotes for countries that themselves quote directly, and direct quotes for countries that themselves quote indirectly.

As already hinted at, in the UK one uses the reverse quote, the number of foreign units that can be bought with one pound, or FC/HC. Some former British-Commonwealth countries (for instance, Australia, New Zealand, and pre-EUR Ireland) do likewise.[3] One reason is that, prior to WW1, the pound was the world's reserve currency and played the role taken over by the dollar after WW2. In addition, until 1967 the GBP was still severely non-decimal—one pound consisted of twenty shilling, each worth twelve pence[4]—while non-pound currencies had gone

[2]Recall from the previous chapter that, until 1972, countries declared an official parity in relation to the USD, say DEM/USD 4. Intervention kept the actual rates between an upper and lower bound expressed, likewise, in DEM/USD.

[3]Canada and South Africa had gone off the pound ages ago, that's why they quote differently.

[4]Recall there also was a dollar (10s), a crown (5s), and a Guinea, worth 21 shillings in the end;

Figure 3.1: **Key exchange rates: pros' notation, dimensions, and nicknames**

Symbol	Currency Pair	dimension	Trading Terminology
USDJPY	US Dollar, in Japanese Yen	JPY/USD	Dollar Yen
USDCHF	US Dollar, in Swiss Franc	CHF/USD	Dollar Swiss, or Swissy
USDCAD	US Dollar, in Canadian Dollar	CAD/USD	Dollar Canada
USDZAR	US Dollar, in South African Rand	ZAR/USD	Dollar Zar or South African Rand
GBPUSD	British Pound, in US Dollar	USD/GBP	Cable
GBPCHF*	British Pound, in Swiss Franc	CHF/GBP	Sterling Swiss
GBPJPY*	British Pound, in Japanese Yen	JPY/GBP	Sterling Yen
AUDUSD	Australian Dollar, in US Dollar	USD/AUD	Aussie Dollar
NZDUSD	New Zealand Dollar, in US Dollar	USD/NZD	New Zealand Dollar or Kiwi
EURUSD	Euro, in US Dollar	USD/EUR	Euro
EURGBP*	Euro, in British Pound	GBP/EUR	Euro Sterling
EURJPY*	Euro, in Japanese Yen	JPY/EUR	Euro Yen
EURCHF*	Euro, in Swiss Franc	CHF/EUR	Euro Swiss
CHFJPY*	Swiss Franc, in Japanese Yen	JPY/CHF	Swiss Yen
GLDUSD	Gold, in US Dollar per troy ounce	USD/ozXAU	Gold
SLVUSD	Silver, in US Dollar per troy ounce	USD/ozXAG	Silver

Key *: cross rate, from the US perspective. Most names should be obvious, except perhaps CHF (*Confederatio Helvetica*, Latin for Switzerland—the way a four-language country solves a political conundrum). The ZAR, South-African Rand, is not to be confused with SAR, Saudi Riyal. GLD and SLV are onorthodox: the official codes as used by e.g. Swift are XAU and XAG, with X signalling a non-standard currency (like also the CFA franc and the Ecu of old), and the Latin *Aurum* and *Argentum*. "Cable" for USDGBP refers to the fact that it is about bank-account money, with payment instructions wired by telegram cable rather than sent by surface mail. There *has* been a time when wiring was cutting-edge technology.

decimal ages ago. It is much easier to multiply or divide by a decimal number, say FC/GBP 0.79208, than with a number like £1/s5/d3 (one pound, five shillings, three pence). So both Brits and non-Brits preferred to talk FC units per pound.

A third (and more recent) class of people using the indirect quote are the Eurolanders, who always quote rates like USD/EUR or JPY/EUR even though they traditionally quoted directly (like DEM/USD). Cynics conjecture that the Europeans may have coveted the reserve-currency status associated with an indirect quote. Another possible reason is that, initially, the Euro was foreign to all existing currencies. For example, to Germans the Euro was introduced as worth two DEM, so they would think it quite naturally to introduce it to Americans and Japanese as worth 1.20 USD or 110 JPY. When, eventually, the Euro had become the home currency, the habit simply stuck.

Example 3.3

Have a look at Figure 3.1, showing the most important rates in the way they are always quoted by pros. The primary rates are in non-US currency except for the

and in Elisabethan times many wages were expressed in marks (13s 4d, i.e. 160d). But by modern times most prices were in £, *s* and *d*.

Figure 3.2: **Sample Spot Exchange Rate Quotes—WSJ Europe**

AMERICAS	Per euro	In euros	U.S. dollar	U.S. dollar	**EUROPE**	Per euro	In euros	U.S. dollar	U.S. dollar
Argentino peso-a	3.9628	0.2523	3.0838	0.3243	**Euro zone** euro	1	1	0.7782	1.2851
Brazil real	2.9588	0.3380	2.3025	0.4343	**Czech Rep** koruna-b	28.260	0.0354	21.992	0.0455
Canada dollar	1.438	0.7073	1.1002	0.9089	**Denmark** krone	7.4576	0.1341	5.8034	0.1723
Chile peso	683.07	0.001464	531.55	0.001881	**Hungary** forint	262.82	0.003805	204.52	0.004890
Columbia peso	3186.28	0.0003138	2479.50	0.0004033	**Malta** lira	0.4294	2.3288	0.3342	2.9926
Ecuador US dollar-f	1.2850	0.7782	1	1	**Norway** krone	7.800	0.1282	6.0698	0.1648
Mexico peso-a	14.5307	0.0688	11.3075	0.0884	**Poland** zloty	3.9369	0.2540	3.0637	0.3264
Peru sol	4.2368	0.2360	3.2970	0.3033	**Russia** ruble-d	34.669	0.02884	26.979	0.03707
Uruguay peso-e	30.841	0.0324	24.000	0.0417	**Slovak Rep** koruna	37.7856	0.02647	29.4040	0.03401
U.S. Dollar	1.2850	0.7782	1	1	**Sweden** krona	9.2662	0.1079	7.2108	0.1387
Venezuala bolivar	2759.39	0.000362	2147.30	0.000466	**Switzerland** franc	1.5604	0.6109	1.2103	0.8262
...	...	...	...	...	...	...	...	...	...

a—floating rate b—commercial rate c—government rate d—Russian Central Bank rate f—Special Drawing Rights from the International Monetary Fund ; based on exchange rates for U.S., British and Japanese currencies.
Note: Based on trading among banks in amounts of $1million and more, as quoted by Reuters

Key: The *Wall Street Journal Europe* sensibly shows both the natural and indirect quotes.

GBP, NZD and AUD or for the EUR; you know why. Cross rates for the EUR are in non-EUR currency, and likewise for the GBP.

Example 3.4

Look at the *Wall Street Journal Europe* excerpt in Figure 3.2, conveniently showing both quotes; the value in USD or EUR of one unit of the third ("foreign") foreign curency, and the value of one USD or EUR in units of that third ("foreign") currency. The natural quote for Americans or Europeans would be the first one, but U.S. traders and Eurolanders may use the other quote. Take a minute to look at Figure **??** and see if you understand the exchange rates as quoted.

Q1. What is the dollar equivalent of one Euro, according to the quotes in *The Wall Street Journal*?
A1. If your answer is USD 1.285, you are correct.

Q2. Determine the amount of Peruvian sol per EUR.
A2. If you answered 4.2368 sol per EUR, you are right.

3.1.4 Bid and Ask Rates

When you deal with foreign currency, you will discover that you pay a higher price at the time of purchase than when you sell one currency for another. For example, for dollar-rouble deals the currency booth in your hotel will quote two numbers, say RUB/USD 35-36. This means that if you sell USD for RUB, you receive RUB 35, while if you wish to buy USD you will have to pay RUB 36. The rate at which the bank will buy a currency from you is called the *bid* rate: they bid (*i.e.* they announce that they are willing to pay) 35 per dollar; and the rate at which the bank will sell a currency to you is the *ask* rate (they ask 36 per dollar). It is, initially, safer not to think about the meaning of bidding and asking because the words refer to the bank's view, not yours. Just remember that *you* buy at the bank's ask rate, and *you* sell at the bank's bid rate. The bid is the lower quote, and ask is the higher one. The Ask comes higher in the alphabet—use any trick that works, until you get used to it.

Indeed, if exchange rates are being quoted with the currency of interest—the currency you are buying or selling—in the denominator, then the ask rate will be higher than the bid rate. Obviously, it could not be the other way around: with a bid rate above the ask rate you would be able to make huge risk-free profits by buying at a the ask and immediately reselling at the assumedly high bid. No bank will allow you to buy low and then immediately resell at a profit without taking any risk, because your sure gains would obviously mean sure losses for the bank. In theory, there could still be room for a situation "bid rate = ask rate" (which offers no such arbitrage opportunities). Yet, the real-world situation is invariably "bid rate < ask rate": banks want to make some money from foreign-currency transactions.

Another way to think of this difference between the ask and the bid rate in fact is that the difference contains the bank's commission for exchanging currencies. The difference between the buying and selling rates is called the *spread*, and you can think of the bank's implicit commission as being equal to half the spread. The following example explains why the commission is half of the spread rather than the spread itself.

Example 3.5

Suppose that you can buy CAD at RUB/CAD 38.6, and sell at RUB/CAD 38.0. With these rates, you can think of a purchase as occurring at the midpoint rate (RUB/CAD 38.3), grossed up with a commission of 0.30. Likewise, a sale can be thought of as a sale at the midpoint, 38.3, from which the bank withholds a commission of 0.30. Thus, the equivalent commission per one-way transaction is the difference between the bid (or ask) and the midpoint rate, that is, half the spread. (The spread itself would be the cost of a round-trip deal—buy and then sell).

To get an idea of whether your house bank charges a low commission, you can ask for a two-way quote to see if the spread is small. If this is the case, you probably do not have to check with other banks. However for large transactions, you should also compare the spot quotes given by different banks. (This will be examined further in Sections 3.3 and 3.3.3.) We discuss the determinants of spreads later, after we have described the market microstructure.

3.1.5 Primary rates *v* cross rates

As of 1945 and until well into the 1980s, all exchange rates in the wholesale segment were against the USD. They were and are called *primary rates*, while any rate not

Figure 3.3: **Cross rates as in the Wall Street Journal Europe**

values of one Euro

Cross rates U.S. dollar and euro foreign-exchange rates in global trading

	USD	GBP	CHF	SEK	RUB	NOK	JPY	ILS	EUR	DKK	CDN	AUD	
Australia	1.3253	2.4818	1.0915	0.1838	0.0491	0.2183	0.0118	0.2934	1.7031	0.2284	1.2046	—	
Canada	1.1002	2.0603	0.9061	0.1526	0.0408	0.1813	0.0098	0.2436	1.4038	0.1896	—	0.8302	
Denmark	5.8034	10.867	4.7794	0.8048	0.2151	0.9561	0.0518	1.2847	7.4576	—	5.2748	4.3789	
Euro	*0.7782*	*1.4573*	*0.6409*	*0.1079*	*0.0288*	*0.1282*	*0.0069*	*0.1723*	—	*0.1341*	*0.7073*	*0.5872*	values in euro
Israel	4.5173	8.4592	3.7202	0.6265	0.1674	0.7442	0.0403	—	5.8049	0.7784	4.1058	3.4085	
Japan	112.11	209.93	92.325	15.547	4.1553	18.469	—	24.817	144.06	19.317	101.90	84.589	
Norway	6.0698	11.367	4.9988	0.8418	0.2250	—	0.0541	1.3437	7.8000	1.0459	5.5170	4.5800	
Russia	26.978	50.521	22.218	3.7414	—	4.4448	0.2407	5.9724	34.669	4.6488	24.521	20.357	
Sweden	7.2108	13.503	5.385	—	0.2673	1.1880	0.0643	1.5963	9.2662	1.2425	6.5541	5.4409	
Switzerland	1.2145	2.2739	—	0.1684	0.0405	0.2000	0.108	0.2688	1.5604	0.2092	1.1037	0.9162	
U.K.	0.5340	—	0.4398	0.0741	0.0198	0.0880	0.0048	0.1182	0.6862	0.0920	0.4854	0.4029	
U.S.	—	1.8726	82.236	0.1387	0.0371	0.1648	0.0089	0.2214	1.2850	0.1723	0.9089	0.7546	

involving the USD would be called a *cross rate* and would traditionally be regarded as just implied by the primary rates. You find an example for midpoint rates in Figure 3.3. The primary rates are in the first column (FC/USD) or the bottom line (USD/FC).[5] The rest of the table is obtained by division or multiplication: GBP/EUR = GBP/USD × USD/EUR, for example. Each of the resulting new rows or columns is a set of quotes in HC/FC (row) or FC/HC (column). With 12 currencies you have 144 entries, of which 12 are on the information-free diagonal, and half of the remaining 132 are just the inverses of the others.

We have a whole section on the relation between primary and cross rates in the presence of spreads, so at this stage we just consider why, among pros, there were just primary rates, until the 1980s. There were several reasons:

- Official parities were against the USD; there was no official parity (in the sense of being defended by any central bank) for rates against other currencies.

- The USD market had the lowest spreads, so all real-world transactions would effectively be done via the dollar anyway. That is, pounds were converted into marks by buying dollars first and then exchanging these for marks, for example, because that was the cheapest way to do so (see below). The cross-rate would just be the rate implied by the two primary rates used in the transaction.

- In pre-electronic days it would be quite laborious to keep track of, say, a 30×30 matrix of cross rates with 435 distinct meaningful entries, making sure all cross rates are consistent with the primary ones all the time. So rather than quoting

[5] Many newspapers give currency j the j-th row and the j-th column instead of the $(N-j)$-th row and the j-th column, but the lay-out is not crucial. The orientation of the empty diagonal (or the unit diagonal, as other tables might show it) is the sign to watch.

cross rates all the time, banks just showed primary quotes and then computed cross rates if and when needed.

By the 1980s desktop computers were around everywhere and, for many pairs of "big" currencies the volume of cross transactions had become large enough to make direct cross exchanges competitive compared to exchanges via the USD. Official exchange rates were gone in many cases, or in the ERM case had become multilateral. So we now see explicit quotes for some of the cross rates. Look at Figure 3.2 to see what rates have active multilateral electronic markets—a good indication of there being a reasonable volume. Note also that for some new EU members the market against the EUR works well while the market against the USD lacks liquidity; that is, for these countries the rate against the euro is economically the key one even though Americans would regard it as just a cross rate.

3.1.6 Inverting Exchange Rates in the Presence of Spreads

The next issue is how a pair of quotes for one currency can be translated into a pair of quotes for a different currency. The rule is that the inverse of a bid quote is an ask quote, and vice versa. To conceptualize this, consider the following illustration.

Example 3.6

An Indian investor wants to convert her CAD into USD and contacts her house bank, Standard Chartered. Being neither American nor Canadian, the bank has no natural preference for either currency and might quote the exchange rate as either USD/CAD or CAD/USD. The Indian bank would make sure that its potential quotes are perfectly compatible. If it quotes from a Canadian viewpoint, the bank gives a CAD/USD quote (which says how many CAD the investor must pay for one USD—for instance, CAD/USD 1.5). If it uses the US perspective, the bank gives a USD/CAD quote, which says how many USD the US investor gets for one CAD, 0.66667.

The bank's alternative ways of quoting will be fully compatible if

$$S_{bid,t}^{CAD/\mathrm{USD}} = \frac{1}{S_{ask,t}^{\mathrm{USD}/CAD}} \tag{3.3}$$

$$S_{ask,t}^{CAD/\mathrm{USD}} = \frac{1}{S_{bid,t}^{\mathrm{USD}/CAD}} \tag{3.4}$$

To fully understand this, recall that what looks like buying (at the ask) to a US resident looks like selling to a Canadian—at the Canadian's bid. Alternatively, recall that the ask is the higher of the two quotes. But if you invert two numbers, the inverse of the larger number will, of course, be smaller than the inverse of the smaller number. Because the inverse of a larger number is a smaller number, the inverse ask must become the bid, and vice versa.

Example 3.7

Suppose that you read the following quote on the Reuters screen: USD/CAD 1.000-1.005.

Q1. What is the bank's buying and selling rate for CAD?
A1. The bank's buying rate for CAD is USD 1.000 and its selling rate is USD 1.005; that is, *you* sell CAD at USD 1.000 and buy at 1.005.

Q2. What, therefore, are the bank's buying and selling rates for USD (in CAD)?
A2. The bank's buying rate or bid for USD is 1/1.005 = CAD/USD 0.995025 (probably rounded to 0.9950) and the selling rate or ask is 1/1.000 = 1.000; that is, wearing your Canadian hat, *you* sell USD at CAD 0.9950 and buy at 1.000.

One corollary is that in countries like the UK, where the reverse or indirect quote is used, the rate relevant when you buy is the lower of the two, while the higher quote is the relevant rate when you sell. Thus, it is important to be aware of what the foreign currency is, and what convention is being used for quoting the exchange rate. Again, it is always easier and more convenient to have the foreign currency in the denominator. That way the usual logic will work: banks buy low and sell high.

3.2 Major Markets for Foreign Exchange

In this section, we describe the size and structure of the exchange market and the type of transactions one can make in this market.

3.2.1 How Exchange Markets Work

The foreign exchange market is not an organized market. Stock markets or futures markets are: they have fixed opening hours, a more or less centralized mechanism to match supply and demand, standardized contracts, an official publication channel for data on volumes and prices, and a specific location or one designated group of computers running everything. In contrast, the exchange market consists of a wholesale tier, which is an informal network of about 500 banks and currency brokerages that deal with each other and with large corporations, and a retail tier, where you and I buy and sell foreign exchange. At any point in time, wholesale exchange markets on at least one continent are active, so that the world-wide exchange market is open twenty-four hours a day (see Figure 3.4). Until the mid-90s, most interbank dealing was done over the telephone; most conversations were taperecorded, and later confirmed by mail, telex, or fax. Reuters—which was already omnipresent with its information screens—and EBS[6] have now built computer networks which

[6] EBS (Electronic Broking Services) was created by a partnership of the world's largest foreign exchange market-making banks. Approximately USD 125 billion in spot foreign exchange transac-

Figure 3.4: **Trader activity over the day**

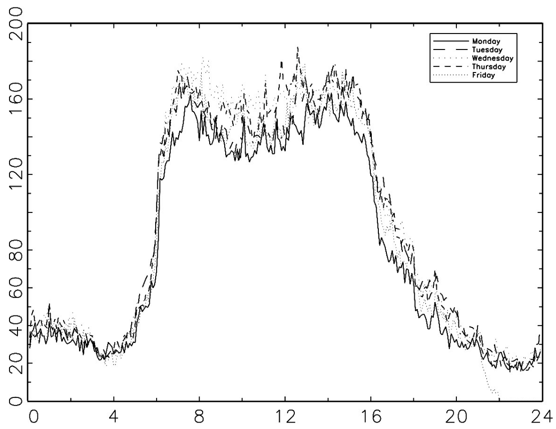

Key Graph courtesy of Luc Bauwens, Université Catholique de Louvain. The graph shows, per 5-minute interval over 24 hours, the evolution of the average number of indicative quotes entered into the Reuters FX/FX pages. Time is GMT in summer, GMT+1 in winter; that is, European time is t+2 hrs, London +1, NY –4 hrs; Sydney and Tokyo time are at t+10 and t+9 hrs, respectively. Below I describe working days as 8:00-17:00, but many a trader starts earlier and/or works later. At 0:00, when the morning shift in Sydney has been up and running for about 2 hrs and Tokyo for 1 hr, Hong Kong starts up, to be followed by Singapore in 1 hr and Bahrain in 3. Between 6:00 and 8:00 the Far East bows out but Western Europe takes over—first the continent (6:00 GMT), then London (7:00) ; activity soars. A minor dip follows around the European noon but activity recovers again in the afternoon, peaking when New York takes over (12:00) and Europeans close their positions (15:00 on the continent, 16:00 London). NY does less and less as time passes. By 22:00 Sydney is starting up, and Tokyo is preparing breakfast.

allow direct trading and that now largely replace the phone market. The way the computer systems are used depends on the role the bank wants to play. We make a distinction between deals via (*i*) market makers, (*ii*) auction platforms, or (*iii*) brokers.

Market Making

Many players in the wholesale market act as market makers. If a market-making credit agreement between two banks has been signed, either party undertakes to provide a *two-way* quote (bid and ask) when solicited by the other party, without even knowing whether that other party intends to buy, or rather sell. Such a quote is *binding*: market makers undertake to effectively buy or sell at the price that was

tion, 700,000 oz in gold and 7 million oz in silver is traded every day over the EBS Spot Dealing System. It was created in 1990 to challenge Reuters' threatened monopoly in interbank spot foreign exchange and provide effective competition. ICAP Plc, the world's largest broker of transactions between banks, agreed in 2006 to buy EBS.

Figure 3.5: **A Reuters conversation and an EBS Broking Window**

A *Reuters* conversation

From: GENP
Hi: EUR/USD in 5 pse?
Hi 25 27 +
Mine 5 at 27 val 5/9 +
Tks $ to Citi Bibi

An EBS broking screen

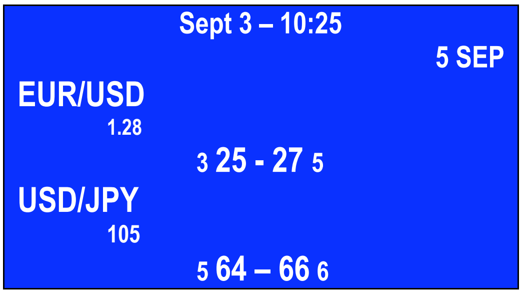

Key In the Reuters conversation window, GENP is an abbreviated name (Jenpi, Jean-Pierre); he asks for a quote for EUR in USD for quantity 5m (dollar); pse is GENPs code for please. The counterparty answers by keying in the small numbers, and Jenpi replies he buys 5 (million) at the ask, 27, for value date Sept 5. The counterpart closes with "Thanks, I'll send the dollars to your correspondent, CitiBank. Bye bye."

The second picture shows part of an EBS broking screen. On top, the current date and time. Next line: the spot delivery date, Sept 5. For two currencies you then see in small font the big figure (the part of the quote that is usually omitted) and in big font the smallquotes: bid and ask, each preceded/followed by the quantity available, in millions. Thus, somebody bids 1.2825 for 3 million dollar, another party offers 5 million dollar at 1.2827.

indicated.

Example 3.8

Deutsche may ask Hong Kong & Shanghai for a quote of EUR against USD. HSBC must then provide a bid and an ask without knowing the direction of Deutsche's possible trade; and if Deutsche replies with "I buy 10 million" then HSBC must sell that quantity at the price they quoted.

Of course there are limits to the market makers' commitments to their quotes. First, potential customers should decide almost immediately whether to buy ("mine"), or to sell ("yours"), or not to deal; they cannot invoke a quote made, say, three minutes ago. Second, if the intended transaction exceeds a mutually-agreed level, laid down in the prior credit agreement—say USD 25m—market makers can refuse. For larger transactions, the trader asking for a quote should reveal immediately what the size of the transaction will be. Third, the credit agreement also provides a limit to the total amount of open contracts that can be outstanding between the two banks at any moment;[7] if the limit is reached, no more deals are allowed.

[7]Exchange transactions are settled with a delay of at least two days, so each contract remains outstanding at least two days; many live much longer. See Section 3.2.3.

Transactions via binding two-way quotes are typically concluded on computers, via kind of chatting windows (more grandly called "conversations"). Bank A's trader X clicks his conversation window with trader Y at bank B—there may be up to 64 such windows open at any given point of time—and might type in, for instance, `PLS EUR/USD`, meaning "please provide a quote for the EUR, in USD". Player A can also mention the quantity, in millions. The millions are omitted; that is, 5 means five million; and the quantity bears on the currency in the denominator, traditionally the USD or the GBP. B's trader may answer, for instance, `13-16` meaning that (the last two digits of) her bid and ask are 13 and 16. (Traders never waste time by mentioning the leading numbers: everybody knows what these are. Only the "small" numbers are mentioned.) The first party can let the offer lapse; if not, he answers `MINE` or `YOURS`, mentions the quantity if not already indicated, and hits the `SEND` key. The deal's done, and both traders now pass on the information to their "back office", which enters the data into the information systems. The back offices will also check with each other to see whether the inputs match; with the logs of the conversations, disputes are of course far less likely than before, when everything went by phone and when traders handed down hand-scribbled "tickets" to the accountants who then checked with each other via telexes. Voice deals still exist, but they are getting rarer.

Implications of Market Making for the Size of the Bid-ask Spread and the Maximum Order Size

Normally, the lower the volume in a particular market, the higher the spread. Also, during holidays, weekends, or lunch times, spreads widen. Spreads are also higher in periods of uncertainty, including the open and close of the market every day. Maximum order quantities for normal quotes follow a similar pattern: a market maker is prepared to handle large lots if the market is liquid (thick) or the volatility low.

All these phenomena are explained by the risk of market making. Notably, if a customer has "hit" a market maker, the latter normally wants to get rid of that new position quickly. But in a thin or volatile market, the price may already have moved against the market maker before he or she was able to close out. Thus, in a thin or volatile market, the market maker wants a bigger commission as compensation for the risk, and puts a lower cap on the size of the deals that can be executed at this spread. For the same reason, quotes for an unusually large position are wide too: getting rid of a very large amount takes more time, and during that time anything could happen. In the retail end of the market, in contrast, the spread increases for smaller transactions. This is because 100 small transactions, each for USD 100,000, cost more time and effort than one big transaction of USD 10m.

For high-volume currencies like the USD/EUR, the difference between one market-maker's own bid and ask is often as low as three basis points (in a quote of four or

Table 3.1: **Order limits and spreads for various rates, semi-professional**

Ticker	Size of 1.0 lot	Instant Execution	Spread	Limit & Stop levels	March 9, 2007 rate (in pips)	spread, %^2
EURUSD	EUR 100,000	up to 10M	2 pips	2 pips	13,115	1.5
GBPUSD	GBP 100,000	up to 10M	3 pips	3 pips	19,319	1.6
EURCHF	EUR 100,000	up to 5M	3 pips	3 pips	16,163	1.9
EURJPY	EUR 100,000	up to 10M	3 pips	3 pips	15,489	1.9
USDJPY	USD 100,000	up to 10M	3 pips	3 pips	11,810	2.5
GBPCHF	GBP 100,000	up to 5M	7 pips	7 pips	23,810	2.9
EURGBP	EUR 100,000	up to 5M	2 pips	2 pips	6,788	2.9
GBPJPY	GBP 100,000	up to 5M	7 pips	7 pips	22,817	3.1
USDCHF	USD 100,000	up to 10M	4 pips	4 pips	12,325	3.2
USDCAD	USD 100,000	up to 5M	4 pips	4 pips	11,735	3.4
AUDUSD	AUD 100,000	up to 5M	3 pips	3 pips	7,802	3.8
CHFJPY	CHF 100,000	up to 5M	4 pips	4 pips	9,583	4.2
EURCAD	EUR 100,000	up to 3M	8 pips	8 pips	15,389	5.2
NZDUSD	NZD 100,000	up to 2M	5 pips	5 pips	9,583	5.2
USDSGD	USD 100,000	up to 1M	8 pips	8 pips	15,267	5.2
EURAUD	EUR 100,000	up to 5M	10 pips	10 pips	16,810	5.9

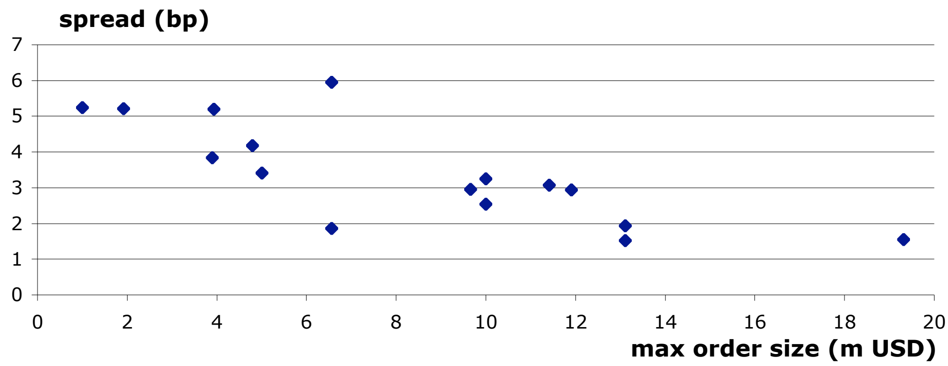

Key The table shows conditions for various currencies from a particular internet broker. The min and max quantities are not interbank, but still aiming at semi-professionals or perhaps day traders rather than pop&mom investors, the hardcore retail. The spread and the tick size for limit and stop levels are likewise wider than interbank. Do note how the spread vary depending on liquidity and the level of the rate, and how the max order size (imperfectly) relates to spread (graph). Source: http://www.alpari.co.uk/en/cspec/ for columns 1-5; WSJ Europe March 12, 2007 for column 6; spread in bp has been added. Data have been re-arranged by increasing relative spread. For the graph the order sizes have been converted from reference currency (the FC in the quote) to USD.

five digits, like 1.2345 or 0.9876), and the difference between the best bid (across all market makers) and the best ask (also across all market makers) may be just two or one or, occasionally, zero basis points. See Section 3.3.2 for more information on quoting behavior.

Table 3.1 shows the minimum and maximum amounts quoted by an internet

Figure 3.6: **A Panel of Reuters Broking Windows**

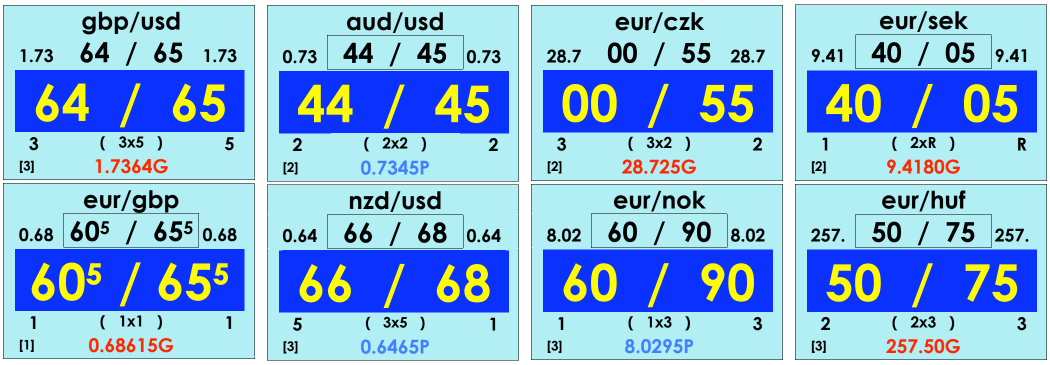

Key The entries should be obvious, by now, except the bottom line, which shows the last trade (quantity and price).

dealer;[8] they are smaller than interbank (and spreads are bigger than interbank), but you can still notice how the maximum amounts and the spreads relate to each other, presumably both reflecting liquidity and volatility.

Auctioning off Through a Broking System

All the above was about market making. Beside these purely bilateral deals—the successors to bilateral phone conversations—there nowadays are increasingly many semi-multilateral deals. If a trader actively wants to buy, or sell, she may enter a limit order into EBS' or Reuters' limit-order book rather than calling a number of market makers or waiting until someone else calls her. This is comparable to you offering, say, a used car for sale on eBay rather than calling various car dealers or posting a sign on your door and then waiting until someone rings your bell. For instance, bank A may have EUR 30m for sale and want at least USD/EUR 1.3007 for them—an ask price. The bank posts this info, for instance, on Reuters' "3000" system. Reuters' window, at any moment, then shows the best bid across all "buy" limit orders, and the best ask among all "sell" limit orders outstanding at that moment. For instance, on Reuters' 3000 screen a line `EUR/USD 10-11 3×R` means that the highest bid posted at that very moment is 10, the lowest ask 11, and that the quantities for these limit orders are, respectively, 3 and "a number exceeding

[8] "Size of one lot" shows the minimum, which is clearly targeting players out of the interbank league (where the lot size is 1m) but still above the micro-investor's league. "Instant execution" is the maximum amount you can buy or sell at the trader's regular quotes.

Table 3.2: **EBS *v* Reuters D2: who leads, who follows, who fails**

	EBS	Reut		EBS	Reut		EBS	Reut
primary			*primary (cont'd)*			*Cross against EUR*		
EUR/USD	+	+/–	USD/RON		–	EUR/AUD		–
AUD/USD	+/–	+	USD/RUB		+	EUR/CAD		–
GBP/USD	+/–	+	USD/SAR		+	EUR/CHF	+	+/–
NZD/USD		+	USD/SEK		–	EUR/CZK	–	+
USD/CAD	+/–	+	USD/SGD	–	+	EUR/DKK	–	+
USD/CHF	+	+/–	USD/THB		+	EUR/GBP	+/–	+
USD/CZK		–	USD/TRY		+	EUR/HUF	–	+
USD/DKK		–	USD/ZAR		+	EUR/ISK		+
USD/HKG	–	+				EUR/JPY	+	+/–
USD/HUF		–	*Other cross rates*			EUR/NOK	–	+
USD/ILS		+	AUD/JPY		–	EUR/NZD		+/–
USD/INR		+	AUD/NZD		+/–	EUR/PLN	–	+
USD/ISK		–	CHF/JPY	–		EUR/RON		–
USD/JPY	+	+/–	GBP/JPY		–	EUR/SEK	–	+
USD/MXN	–	+				EUR/SKK		+
USD/NOK		–				EUR/TRY		–
USD/PLN	–	+				EUR/ZAR		+/–

+: primary liquidity source; +/–: supported, but liquidity not good or not stable; –: supported but not used in practice. Rates are expressed following the "name" convention, not the dimensions. **Source:** www.londonfx.c.uk/autobrok.html, feb 2007

50" (=R).[9] You see the EBS counterpart of Reuters 3000 in Figure 3.5. Any party interested in one of these offers can then click on the quote they like (either the bid or ask) and specify the quantity taken. Or another bank may enter a limit order that is automatically matched, wholly or partly, with an already outstanding limit order. Reuters' computer then informs the IT systems of both banks of the transactions that were concluded so that no more human intervention with "tickets" and telexes and faxes is needed (*straight through processing*, STP).

The decision by an FX trader whether to use EBS or Reuters Dealing 3000 (aka D2) is driven largely by currency pair. In practice, EBS is used mainly for EUR/USD, USD/JPY, EUR/JPY, USD/CHF and EUR/CHF, and Reuters D2 is used for all other interbank currency pairs. Have a look at Table 3.2 to see who leads where. In these multilateral electronic dealing systems, the spread for EUR/USD is typically one pip, that is, one hundredth of a USD cent. (Online currency brokers targeting private investors typically offer a 2-pip spread; just feed "foreign exchange" into your web search engine to find these brokers.) For other exchange rates spreads are often

[9] The quotes are, again, "small numbers" and the quantities mean millions of dollars. Remember also that, for traders, EUR/USD means "the value of the Euro in dollars".

Figure 3.7: **The Bank of International Settlements (BIS)**

The Bank for International Settlements (BIS) is commonly described as the bank of the Central Banks. It was first set up after WW1 to act as a payent agent distributing the German and Austrian war reparation payments. After WW2 it ran the European Payment Union (EPU), serving as a netting institute for payments among EPU members. By netting the international payments, the volume of actual payments was reduced, which alleviated the problems of dollar shortages in the first years after the war. Currently, the BIS still is the bank of the central bankers: all central banks have accounts there, in various currencies, and can route their payments to each other via the BIS. But nowadays the BIS mainly serves as a talking club for central bankers and regulators. One of its missions is to gather data on exchange markets, euro- and OTC-markets, new financial instruments, bank lending to sovereign borrowers, and so on. Another mission is to provide a forum where regulators coordinate the *capital adequacy* rules that they impose on financial institutions. The Basel-1 rules covered credit risk—in a crude way, perhaps, but it was a useful first step; the recent Basel-2 rules refine Basel-1 and add market-price risks, see the chapter on Value at Risk.

wider.

Note that the advent of these multilateral systems has made the market somewhat more like an organized market: there is centralization of buy and sell orders into one matching mechanism, there are membership rules (not anyone can log on into the program), rules about orders, etc. But the exchange market is still fully private, whereas many exchanges are semi-official institutions that are heavily regulated and need, at least, a license.

Brokers

A last way of shopping around in foreign exchange markets is through currency brokers. In the telephone-market days, brokers used to do the middleperson stuff that nowadays is handled via limit-order books: on behalf of a bank or company, the broker would call many market makers and identify the best counterpart. Roughly half of the transaction volume in the exchange market used to occur through brokers. Nowadays, brokers are mainly used for unusually large transactions, or "structured" deals involving, say, options next to spot and/or forward; for bread-and-butter deals their role is much reduced.

3.2.2 Markets by Location and by Currency

Every three years, in April, the Bank of International Settlements (see box in Figure 3.7) makes a survey of the over-the-counter markets, including forex. At the latest count, April 2007, the daily volume of trading on the exchange market and its satellites—futures, options, and swaps—was estimated at more than USD 3.2 trillion. This is over 45 times the daily volume of international trade in goods and services, 80 times the US' daily GDP, 230 times Japan's GDP, and 400 times Germany's GDP, and

Figure 3.8: **Forex turnover, daily, USDb, and market shares of currency pairs**

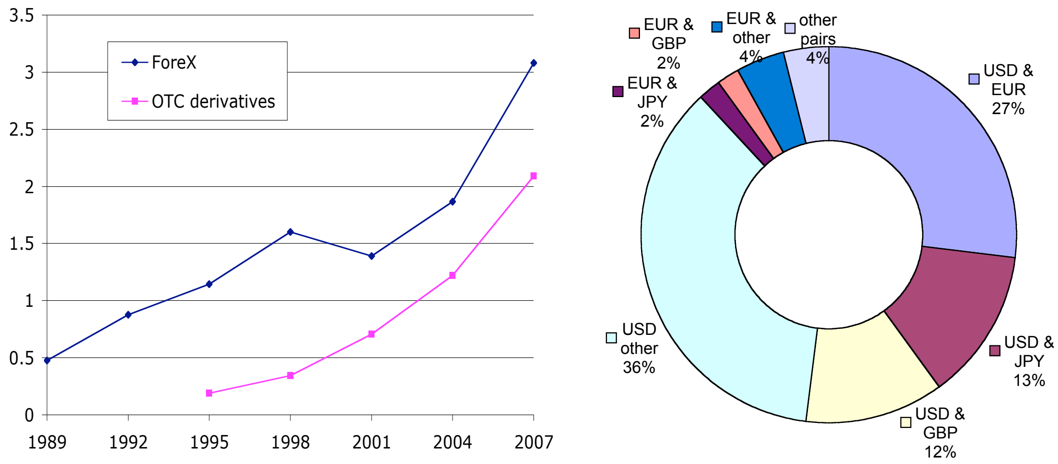

Source BIS, *Triennal Central Bank Survey of Foreign Exchange and Derivatives Market Activity in April 2007*, Preliminary global results, September 2007.

7500 times the world's official development-aid budget.[10] The major markets were, in order of importance, London, New York, Tokyo, Frankfurt (the European Central Bank's home base). London leads clearly, easily beating even New York, Tokyo, and Singapore taken together, and still increasing its market share. Frankfurt is a fast riser but from a low base.

The most important markets, per currency, are the USD/EUR and the USD/JPY markets; together they represent almost half of the world trading volume. Add in the GBP, and the transactions involving just the top four moneys represent two thirds of all business. The USD still leads: in 88 % of transactions it takes one of the sides (down from 90 in 2004), while the EUR is one of the two currencies in less than 40% (up from 35) of that volume—and the bulk of that is USD/EUR trade.

3.2.3 Markets by Delivery Date

The exchange market consists of two core segments—the spot exchange market and the forward exchange market.

The **spot** market is the exchange market for quasi-immediate payment (in home currency) and delivery (of foreign currency). For most of this text we shall denote

[10]All data are from the CIA Factbook—see Google. Trade and aid: 2004; GDP: early 2007 estimates for 2006.

Table 3.3: **Market shares. %, for foreign exchange trading**

	UK	US	Japan	Singa-pore	Other
1998	32.5	17.9	6.9	7.1	35.6
2001	31.2	15.7	9.1	6.2	37.8
2004	31.3	19.2	8.3	5.2	36.0
2007	34.1	16.6	6.0	5.8	37.5

Source BIS, *Triennal Central Bank Survey of Foreign Exchange and Derivatives Market Activity in April 2007*, Preliminary global results, September 2007.

this spot rate by S_t, with t referring to current time. In practice, "quasi-immediate" means "right now" only when you buy or sell notes or coins. (This section of the market is marginal.) For electronic money (that is, money that will be at your disposal in some bank account), delivery is in two working days for most currencies ("$t + 2$"), and one day between Canada and the US or between Mexico and the US ("$t + 1$"). Thus, if you buy AUD 2m today, at AUD/EUR 2, the AUD 2m will be in your account two working days from now, and the EUR 1m will likewise be in the counterpart's account two days from now. The two-day delay is largely a tradition from the past, when accounts were kept by hand. The hour of settlement depends on the country, but tends to be close to noon. Thus, the EUR side of a EUR/USD transaction is settled in Europe about six hours before the USD leg of the seal is settled in NY.[11]

The **forward** market is the exchange market for payment and delivery of foreign currency at some future date, say, three months from now. For example, supposing today is January 3, you could ask your bank to quote you an exchange rate to sell dollars for pounds for a date in March, say March 5, and the transaction would be settled on that date in March, at the rate agreed upon on January 3 (irrespective of the spot rate prevailing on March 5). The forward market, in fact, consists of as many subsegments as there are delivery dates, and each subsegment has its own price. We shall denote this forward rate by $F_{t,T}$, with T referring to the future delivery date. (Forward rates and their uses will be discussed in great detail in Chapters 4 and 5.)

The most active forward markets are for 30, 90, 180, 270, and 360 days, but nowadays bankers routinely quote rates up to ten years forward, and occasionally even beyond ten years. Note that months are indicated as thirty days. In principle, 30-day contract is settled one month later than a spot contract, and a 180-day for-

[11]This leads to the risk that, in between the two settlement times, one party may file for bankruptcy or be declared bankrupt. This is called "Herstatt risk", after a small German bank that pulled off this feat on June 26, 1974. Nowadays, regulators close down banks *outside* working hours.

ward contract is settled six months later than a spot contract—each time including the two-day initial-delay convention.[12]

Example 3.9

A 180-day contract signed on March 2 works as follows. Assuming that March 4 is a working day, spot settlement would have been on March 4. For a 180-day forward deal, the settlement date would be moved by six months to, in principle, September 4, or the first working day thereafter if that would have been a holiday. The actual number of calendar days is at least (2+)184 days: there are four 31-day months in the March-September window.

The above holds for standard dates, But you can always obtain a price for a "*broken date*" (i.e. a non-standard maturity), too. For instance, on April 20 you can stipulate settlement on November 19 or any other desired date.

Worldwide, spot transactions represent less than 50 percent of the total foreign-exchange market volume. The forward market, together with the closely related swap market (see Chapter 7), make up over 50 percent of the volume. About 3 percent of total trade consists of currency-futures contracts (a variant of forward contracts traded in secondary markets—see Chapter 6) and currency options (see Chapter 8).

After this digression on the meaning of exchange rates and their relation to real quantities, we now return to the operations of the spot exchange market. We want to introduce one of the cornerstones of finance theory, the Law of One Price.

3.3 The Law of One Price for Spot Exchange Quotes

In frictionless markets, two securities that have identical cash flows must have the same price. This is called the Law of One Price. There are two mechanisms that enforce this law. The first one is called *arbitrage* and the second one can be called *shopping around*. We explain these two concepts below.

Suppose that two assets or portfolios with identical cash flows do not have the same price. Then any holder of the overpriced asset could simultaneously sell this asset and buy the cheaper asset instead, thus netting the price difference without taking on any additional risk. If one does not hold the overpriced asset, one could still take advantage of this mispricing by short-selling[13] the overpriced asset and covering this with the purchase of the cheaper security. For example, you sell an overpriced

[12] Further details of settlement rules are provided in Grabbe (1995).

[13] see box in Figure 3.9.

Figure 3.9: **What's shortselling?**

In a shortsale you hope to be able to buy low and sell high, but with the selling preceding the buying, unlike in a long position. Thus, a shortseller hopes to make money from falling prices rather than from rising prices.

In markets with delivery a few working days later, you can always go short for a few hours: sell "naked" in the morning, for instance, and then buy later within the same day so as to be able to deliver n days later.

For longer horizons one needs more. In the case of securities, shortselling then requires borrowing a security for, say, a month and selling it now; at the end of the month you then buy back the number of securities you borrowed and restitute them to the asset lender, including dividends if any were paid out during that period.

For currencies, longer-term shortselling can be done by just borrowing forex and selling it, hoping to be able to buy back the forex (including interest) later at a lower price. If there is a forward market, lastly, going short is even easier: promise to deliver on a future date at a price that is fixed now. If prices have dropped by then, as you hope, you'll be able to close out (buy spot) cheaply and make money on the forward deal.

asset at 1.2135 and buy a perfect substitute at 1.2133, netting 0.0002 per unit right now and no net cash flow at T. Such transactions are called arbitrage. These arbitrage transactions generate an excess supply of the overpriced asset and an excess demand for the underpriced asset, moving the prices of these two assets towards each other. In frictionless markets, this process stops only when the two prices are identical. Note that apart from the arbitrage gain, an arbitrage transaction does not lead to a change in the net position of the arbitrageur; that is, it yields a sure profit without requiring any additional investment.

The second mechanism that enforces the Law of One Price is shopping around. Here, in contrast to arbitrage, investors do intend to make particular changes in their portfolios. Shopping around has to do with the fact that, when choosing between different ways of making given investments, clever investors choose the most advantageous way of doing so. Therefore, when choosing between assets with identical cash flows, investors buy the underpriced assets rather than the more expensive ones. Likewise, when choosing which assets to sell, investors sell the overpriced ones rather than the ones that are relatively cheap. This demand for the underpriced assets and supply of the overpriced ones again leads to a reduction in the difference between the prices of these two securities.

Although the arbitrage and shopping-around mechanisms both tend to enforce the Law of One Price, there are two differences between these mechanisms.

- First, an arbitrage transaction is a round-trip transaction. That is, you buy and sell, thus ending up with the same position with which you started. As arbitrage requires a two-way transaction, its influence stops as soon as the price difference is down to the *sum* of the transactions costs (buying and selling). In contrast, in shopping around one wishes to make a particular transaction, and the issue is which of the two assets is cheaper to trade.[14] As a result, the influence of shopping around can go on as long as the price difference exceeds the *difference* of the two transactions costs.[15]

- Second, arbitrage is a strong force because it does not require any capital and can, therefore, generate enormous volumes. In contrast, shopping around can be a price-equalizing mechanism only if there are investors who wish to make that particular transaction. This exogenously triggered volume, if any, is always finite and may be exhausted before it has actually equalized the prices.

In this section, we apply these arguments to spot rates quoted for the same currencies by different market makers. In a perfect exchange market with zero spreads, arbitrage implies that the rate quoted by bank X must equal the rate quoted by bank Y: there can be only one price for a given currency—otherwise, there is an arbitrage opportunity.

Example 3.10

If Citibank quotes DEM/USD 1.6500, while Morgan Chase quotes DEM/USD 1.6501, both at zero spreads, then

- there is an arbitrage opportunity. You can buy cheap USD from Citibank and immediately sell to Morgan Chase, netting DEM 0.0001 per USD. You will, of course, make as many USD transactions as you can. So will everybody else. The effect of this massive trading is that either Citibank or Morgan Chase, or both, will have to change their quotes so as to stop the rapid accumulation of long or short positions. That is, situations with arbitrage profits are inconsistent with equilibrium, and are eliminated very rapidly.
- there is also a shopping-around pressure. All buyers of USD will buy from Citibank, and all sellers will deal with Morgan Chase.

The only way to avoid such trading imbalances is if both banks quote the same rate.[16]

---

[14] Accordingly, Deardorff (1979) refers to standard arbitrage as *two-way arbitrage* and to shopping-around as *one-way arbitrage*.

[15] Denote by P_U and k_U the price and transaction cost when dealing in the underpriced asset, and denote by P_O and k_O the counterparts for the overpriced asset. The advantage of buying the cheap asset rather than the expensive one remains positive as long as $P_U + k_U < P_O + k_O$; that is, as along as $P_O - P_U > k_U - k_O$. In contrast, the advantage of buying the cheap asset and selling the expensive one remains positive as long as $P_O + k_O - (P_U - k_U <) > 0$; that is, as along as $P_O - P_U > k_U + k_O$: you pay both costs instead of replacing one by another.

[16] This is often put as "by arbitrage, the quotes must be the same," or "arbitrage means that the quotes must be the same." Phrases like this actually mean that to rule out arbitrage opportunities, the quotes must be the same.

Figure 3.10: **Arbitrage and Shopping-around opportunities across market makers**

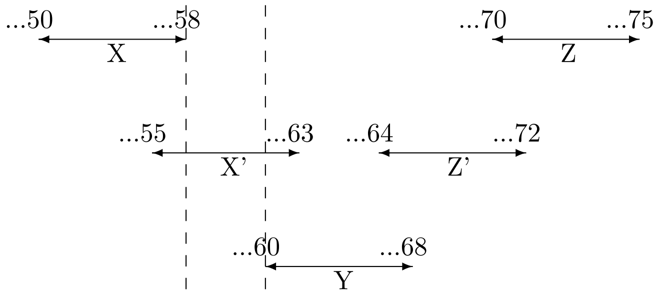

What we now want to figure out is how arbitrage works when there are bid-ask spreads. The point is that, because of arbitrage, the rates cannot be systematically different; and if the quotes do differ temporarily, they cannot differ by too much.

3.3.1 Arbitrage across Competing Market Makers

Suppose bank X quotes you INR/NZD 20.150–20.158 while bank Y quotes INR/NZD 20.160–20.168. If you see such quotes, you can make money easily: just buy NZD from bank X at INR 20.158, immediately resell it to bank Y at INR 20.160, and pocket a profit worth INR 0.002 for each NZD. Note two crucial ingredients: (1) you are not taking any risk, and (2) you are not investing any capital since the purchase is immediately reversed and both transactions are settled on the same day. The fact that you immediately reverse the transaction explains why this is called arbitrage.

If such quotes are found in the exchange market (or elsewhere, for that matter), large trades by a few alert dealers would immediately force prices back into line. The original quotes would not be equilibrium quotes. In equilibrium, the arbitrage argument says that you cannot make money without investing capital and without taking risk. Graphically, any empty space between the two quotes would correspond to an arbitrage profit. Thus, the no-arbitrage condition says that any two banks' quotes should not be separated by empty space; that is, they should overlap by at least one point, like the quotes X' and Y in Figure 3.10.

3.3.2 Shopping Around across Competing Market Makers

Shopping-around activity implies that small differences like those between the pair (X', Y) in Figure 3.10 will not persist for very long. Rather, the two quotes will sometimes be the same, and if at other times one bank is more expensive then this would say very little about what the situation will be five minutes later. To see this, suppose that bank X' quotes INR/NZD 20.55–20.63, while bank Y quotes INR/NZD 20.60–20.68. In such a situation, all buyers of NZD will, of course, prefer to deal with bank X', which has the lower ask rate (20.63 instead of 20.68), while all sellers will now deal with bank Y, which has the better bid rate (20.60 instead of 20.55). It is conceivable that these banks actually want this to happen—for instance if bank X' has an excess of foreign currency (long), and bank Y is short forex and wants to replenish its FC inventory. But we would not expect this to be a long-run phenomenon. It is true that very often a bank may want one type of transaction only, but situations like that must change very rapidly because otherwise that bank's position would become unacceptably large and risky.

Example 3.11

Suppose you see five banks quoting EUR against USD, as follows:

Citibank	USD/EUR	1.3450-52
Bank of America	USD/EUR	1.3450-52
Continental Bank	USD/EUR	1.34<u>51</u>-53
Deutsche Bank	USD/EUR	1.3450-52
Banca da Roma	USD/EUR	1.3449-<u>51</u>

Q. Which bank(s) is (are) keen on buying EUR? keen on selling EUR? not interested in dealing?
A. Continental, with its high bid, is quite attractive to sellers, so this trader clearly wants to buy—for example to fill a short position or because she expects a rise. Roma, in contrast, judging by its low ask, is quite attractive to buyers, so their trader clearly wants to sell—maybe to move an unwanted long position, or in anticipation of a fall in the rate. The others are just twiddling thumbs: as things stand, they are unwilling to match Continental's or Roma's rates, and they hope that things will be better soon.

Q. Why does Continental rise both its bid and its ask, rather than just its bid?
A. Apparently it wants not just to attract sellers but also to scare off buyers. Similarly, Roma not just fancies buyers, but does not want any sellers at all.

Q. If we would look at these banks' quotes every five minutes, do we always expect to see the same pattern, i.e. Continental quoting higher and Roma lower than the majority?
A. Of course not: as soon as their desired positions are reached, they will return to the fold. Thus, the top and bottom positions are picked by a particular bank for only a brief period, and move randomly across the list of banks.

3.3.3 Triangular Arbitrage

Now that we know how exchange rates are quoted and what arbitrage means, let us look at the relationships that exist between spot rates quoted in various currencies.

Figure 3.11: **Triangular arbitrage and triangular shopping around**

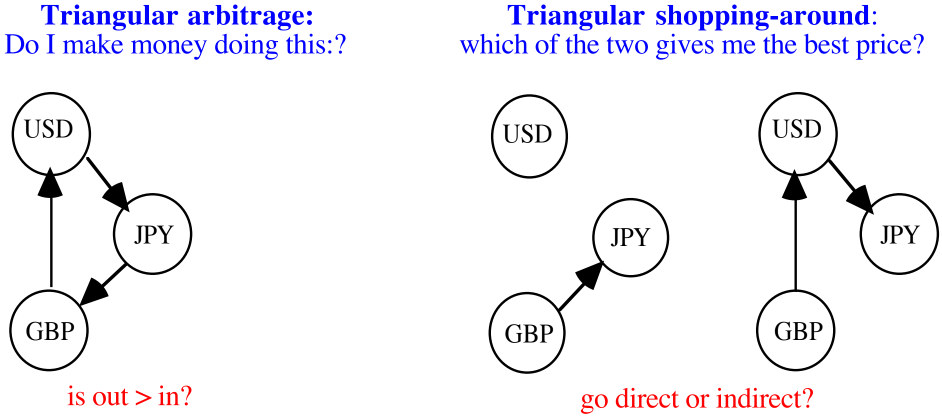

The forces that support these linkages are again arbitrage and shopping around. For our purposes, we can ignore the many market makers: when we talk about bid and ask, we now mean the *market quote*, that is, the best bid across all market makers, and the best ask. The new issue is how these market quotes in various currencies are linked.

- Someone engaging in triangular *arbitrage* tries to make money by sequentially buying and selling various currencies, ending with the original currency. For instance, you could convert AUD into USD, and then immediately convert the USD into GBP and the GBP back into AUD, with the hope of ending up with more AUD than you started out with. The no-arbitrage condition says that you should not make a profit from such activities. Actually, when there are transactions costs or commissions, you are likely to end up with a loss. The potential loss is due to commissions, notably the bid-ask spread. Thus, in this context, arbitrage implies that the set of exchange rates quoted against various base currencies should be such that you cannot make any risk-free instantaneous profits after paying transactions costs.

- *Shopping around* is the search for the best way to achieve a desired conversion. For instance, an Australian investor who wants to buy GBP may buy directly, or may first convert AUD into USD and then convert these USD into GBP. Shopping around implies that the direct AUD/GBP market can survive only if its quotes are no worse than the implied rates from the indirect transaction.

In the case of perfect markets, the regular arbitrage and shopping-around arguments lead to the same conclusion. We illustrate this in the following example.

Example 3.12

Suppose one GBP buys USD 1.5, while one USD buys AUD 1.6; therefore, if we directly convert GBP into Aussies, one GBP should buy $1.5 \times 1.6 = 2.4$ AUD. With this AUD/GBP rate and assuming a zero spread,

- nobody can make a free-lunch profit by any sequence of transactions, and
- everyone is indifferent between direct conversions between two currencies and indirect, triangular transactions.

Below, we see what the implications of arbitrage and shopping around are when there are bid-ask spreads. In order to simplify matters, we shall first show how to compute the implied rates from an indirect route. We shall call these implied rates the *synthetic* rates. Having identified these synthetic rates, we can then invoke the same mechanisms that enforce the Law of One Price as when we studied the relationship between the quotes made by various market makers.

Computing Synthetic Cross-Rates

In general, a synthetic version of a contract is a combination of two or more other transactions that achieves the same objective as the original contract. That is, the combination of the two or more contracts *replicates* the outcome of the original contract. We shall use the notion of replication repeatedly in this textbook. For now, consider a simple spot transaction: a Japanese investor wants to convert JPY into GBP.

- The investor can use the direct market and buy GBP against JPY. We will call this the original contract.
- Alternatively, the investor can first buy USD with JPY, and then immediately exchange the USD for GBP. This is a combination of two contracts. It replicates the original contract since, by combining the two transactions, the investor initially pays JPY, and ultimately ends up with GBP. Thus, this is a synthetic way of achieving the original transaction.

Note that the synthetic contract may be the more efficient way to deal, since the USD market has a lot of volume (or depth) in every country, and therefore has smaller spreads. (This is why the USD is involved in 90 percent of the trades). Let us see how the synthetic JPY/GBP rates can be computed.

Example 3.13

What are the synthetic JPY/GBP rates, bid and ask, if the quotes are JPY/USD 101.07 - 101.20 and USD/GBP 1.3840 -1.3850?

Step 1: multiply or divide? The dimension of the rate we are looking for is JPY/GBP. Because the dimensions of the two quotes given to us are USD/GBP and

Figure 3.12: **A Dealer's Shopping-Around Spreadsheet**

C2 =RtGet("IDN":"EUR=EBS":"BID")

	A	B	C	D	K	L	M	N	O	P	Q	R	S	T	U	V
2	EUR	USD	1,4868	1,4869	EBS	1,4868	1,4869	27,1302	27,1320			EUR/NOK		USD/NOK		
3	EUR	JPY	157,9300	157,9400	RD2002	1,4868	1,4869				RD2002	8,0745	8,0765	5,4304	5,4321	Via EUR
4	EUR	GBP	0,7477	0,7478	Reuters	1,4868	1,4871				Reuters	8,0758	8,0768	5,4303	5,4323	Reuters
5	EUR	CHF	1,6050	1,6053		EUR/GBP		GBP/USD				EUR/DKK		USD/DKK		
6	EUR	AUD	1,6673	1,6683	RD2002	0,7477	0,7478	1,9884	1,9887	RD2002	RD2002	7,4527	7,4530	5,0122	5,0128	Via EUR
7	EUR	CAD	1,4846	1,4853	Via USD	0,7476	0,7478	1,9884	1,9886	Via EUR	Reuters	7,4533	7,4535	5,0123	5,0128	Reuters
8	EUR	CZK	26,062	26,08	Reuters	0,7477	0,7480	1,9884	1,9887	Reuters		EUR/CAD		USD/CAD		
9	EUR	DKK	7,4527	7,4530		EUR/JPY		USD/JPY			Via USD	1,4846	1,4851	0,9985	0,9988	RD2002
10	EUR	EEK	15,645	15,648	EBS	157,93	157,94	106,21	106,22	EBS	Reuters	1,4846	1,4853	0,9986	1,6539	Reuters
11	EUR	HKD	11,5940	11,5960	RD2002	0,00	0,00	106,21	106,23	RD2002		EUR/AUD		AUD/USD		
12	EUR	HUF	259,41	259,61	Via USD	157,91	157,94	106,21	106,23	Via EUR	Via USD	1,6676	1,6681	0,8914	0,8916	RD2002
13	EUR	IDR	13577	13877	Reuters	157,93	157,97	106,20	106,24	Reuters	Reuters	1,6673	1,6683	0,8914	0,8916	Reuters
14	EUR	ISK	96,75	96,82		EUR/CHF		USD/CHF				EUR/NZD		NZD/USD		
15	EUR	LTL	3,4526	3,4528	EBS	1,6050	1,6053	1,0795	1,0796	EBS	Via USD	1,8955	1,8963	0,7841	0,7844	RD2002
16	EUR	LVL	0,6973	0,6983	Via USD	1,6050	1,6053	1,0794	1,0797	Via EUR	Reut/Re	1,8950	1,8968	0,7837	0,7847	Reuters
17	EUR	MYR	4,8086	4,8137	Reuters	1,6049	1,6052	1,0795	1,0798	Reuters		EUR/SGD		USD/SGD		
18	EUR	NOK	8,0745	8,0765		EUR/SEK		USD/SEK			Via USD	2,1074	2,1077	1,4174	1,4175	RD2002
19	EUR	NZD	1,8950	1,8968	RD2002	9,4715	9,4735	6,3700	6,3717	Via EUR				1,4172	1,4177	Reuters
20	EUR	PHP	60,19	60,257	Reuters	9,4726	9,4756	6,3698	6,3708	Reuters	koers in kleine pips		26,5			

Key Courtesy of Paul Goossens, dealer at KBC Brussels. Paul's spreadsheet shows the best quotes from EBS's broking screens, from Reuters Dealing 2002, and the indirect quotes (via USD or EUR). The latter are obviously rounded. Check how the indirect quotes are always wider at one side at least. (With only two pips between the best direct quotes, and with rounding of the synthetic quotes, one side must always seem to match.) The wider quotes labeled Reuters are the indicative, non-binding ones from the `Reuters FX/FX` pages; they mean nothing except that that some banks are willing to quote. See how Paul's sheet gets the EUR/USD quote from EBS into the black&green part of the spreadsheet. Cell 1 is selected; spot the underlying command `=RtGet(''IDN'':''EUR=EBS'':''BID'')` in the `enter function` box above the spreadsheet. From the imported data in the black part, synthetic rates are computed.

JPY/USD, the way to obtain the synthetic rate is to multiply the rates, as follows:

$$[\text{JPY/GBP}\ ] = [\text{JPY/USD}\ ] \times [\text{USD/GBP}\ ] \tag{3.5}$$

Note that on the right-hand side of the equation, the USD in the denominator of the first quote cancels out with the USD in the numerator of the second quote, leaving us with the desired JPY/GBP number.

Step 2: bids or asks? The first quote is the natural quote for a Japanese agent, the second one takes the USD as the base. Consider the synthetic ask (relevant for buying GBP from a JPY position). Starting from JPY we buy USD, so we need the ask; and with the USD we buy GBP, so we again need ask. Thus,

$$\begin{aligned} Synthetic\ S_{t,ask}^{\text{JPY/GBP}} &= S_{t,ask}^{\text{JPY/USD}} \times S_{t,ask}^{\text{USD/GBP}} \\ &= 101.20 \times 1.3850 = 140.16 \end{aligned} \tag{3.6}$$

By a similar argument, we can obtain the rate at which we can synthetically sell GBP into USD and these into JPY:

$$\begin{aligned} Synthetic\ S_{t,bid}^{\text{JPY/GBP}} &= S_{t,bid}^{\text{JPY/USD}} \times S_{t,bid}^{\text{USD/GBP}} \\ &= 101.07 \times 1.3840 = 139.88 \end{aligned} \tag{3.7}$$

This example is the first instance of the Law of the Worst Possible Combination or the Rip-Off Rule. You already know that for any single transaction, the bank gives you the worst rate from your point of view (this is how the bank makes money). It follows that if you make a sequence of transactions, you will inevitably get the worst possible cumulative outcome. This Law of the Worst Possible Combination is the first Fundamental Law of Real-world Capital Markets. In our example, this law works as follows:

- Note that we are computing a product. The synthetic ask rate for the GBP (the higher rate, the one at which you buy) turns out to be the highest possible product of the two exchange rates: we multiply the two high rates, ask times ask. Note finally that, if the purpose is to buy forex, then a high rate is also an unfavorable rate. In short, we buy at the worst rate, the highest possible combined rate.
- We see that, likewise, the synthetic bid rate for the GBP (the lower rate, the one at which you sell) turns out to be the lowest possible product of the two exchange rates: we multiply the two low rates, bid times bid. Note also that, if the purpose is to sell forex, then a low rate is also an unfavorable rate. In short, we sell at the worst rate, the lowest possible combined rate.

Let us look at another example. The data are the same except that the British quotes now are direct not indirect.

Figure 3.13: **Triangular Arbitrage and Shopping-around**

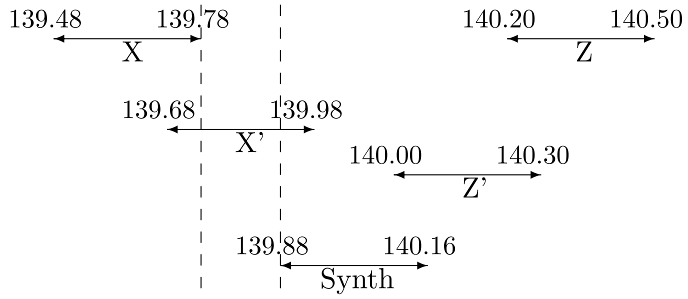

DoItYourself problem 3.1

The JPY/GBP synthetic bid and ask rates, if the quotes are
JPY/USD 101.07 - 101.20 and
GBP/USD 0.72202 - .72254,
are

$$Synth\ S_{t,bid}^{\mathrm{JPY/GBP}} = \frac{S_{t,bid}^{\mathrm{JPY/USD}}}{S_{t,ask}^{\mathrm{GBP/USD}}} = 139.88\,,\ Synth\ S_{t,ask}^{\mathrm{JPY/GBP}} = \frac{S_{t,ask}^{\mathrm{JPY/USD}}}{S_{t,bid}^{\mathrm{GBP/USD}}} = 140.16. \tag{3.8}$$

- Derive this solution from the previous one, invoking our earlier results on inverse rates, equations (3.3) and (3.4).
- Verify that you get the above answer also if you first think of the dimensions and then apply the Law of the Worst Possible Combination.

Triangular Arbitrage with Transactions Costs

Now that we understand synthetic quotes, we can derive bounds imposed by arbitrage and shopping around on quotes in the wholesale market. Just think of the direct quotes as the quotes from bank X, and think of the synthetic quotes as the quotes from bank Y.

- *Arbitrage* then says that the two bid-ask quotes should overlap by at least one point; otherwise, you can buy cheap in the direct market and sell at a profit in the synthetic market or vice versa.
- *Shopping around* implies that if a bank skews its quotes so as to be (very) attractive at (only) one side, then it will attract a lot of business very fast; thus, this skewing cannot be persistent. But when we talk about market

quotes (the best bid, and the best ask, across all market makers) rather than the quotes by an individual dealer, the force is even stronger. Individually, a market maker may very well want to make one of its quotes unappealing for some time, as we saw. But if there are many market makers it would be quite unlikely that, across all market makers, even the best direct quote would still be unappealing against the synthetic one, for that would mean that among all the competing market makers there is not a single one that is interested in that particular type of deal. Thus, instances where a direct quote is dominated by a synthetic one at one side should be rare and short-lived, and the more so the higher the number of market makers.

- The above assumes that the direct market has enough volume. Indeed, with a very thin market, the spread required to make market-making sustainable may be too wide to allow the direct market to compete on both sides with the synthetic market via a heavily-traded vehicle currency (like the USD or the EUR). The volume and depth of the wholesale market for dollars relative to almost any other currency is so large (and the spreads, therefore, so small) that a substantial part of the nondollar transactions are, in fact, still executed by way of the dollar. Direct cross-deals have emerged as of the mid 1980s only, and are still confined to heavy-volume currency pairs.

As a final note, in the retail markets most customers have no direct access to cross rates, and bank clerks occasionally compute cross rates even where the actual transaction could be executed very differently. A Japanese bank, for instance, would post quotes for JPY/GBP and JPY/EUR rates for its retail customers, but typically not for GBP/EUR. Should a retail customer sell EUR and buy GBP, the clerk would actually compute the synthetic rates we just derived, as if the customer first went from EUR to JPY and then to GBP, even if in the bank's trading room the actual conversion may be done directly from EUR into GBP. Unless you have an account with a Euroland or UK bank, or enough clout with your home bank, you would have little choice but to accept the large spread implied by such synthetic rates.

This finishes our tour of the workings of the exchange markets. We continue the chapter with some wise advice on the merits and shortcomings of using exchange rates to translate foreign amounts of money. This brings us to the twin concepts of "PPP" and "real" exchange rates, key issues to understand the relevance of currency risk.

3.4 Translating FC Figures: Nominal rates, PPP rates, and Deviations from PPP

Obviously, when you exchange a FC amount into HC or *vice versa*, you will use the exchange rate relevant at the moment. But actual transactions like this are not the sole conceivable purpose for such a conversion; rather, the purpose may just be *translation*, that is, to have an idea what a FC amount means in a unit that you are more familiar with, the HC. For instance, if a resident of Vanuatu tells you she's making 1m Vatus a month, most people would not have a clue whether they should be impressed or not. In a case like this we don't want to actually exchange any Vatus into our own HC; we'd simply like to translate a FC number into a (to us) more meaningful unit.

The most commonly used solution is to resort to the market exchange rate to make the translation. The result is an improvement on the FC amount in the sense that you know what you would be able to do with this converted amount if you consumed it here, at home. But your objective may be to have a feel for what the FC amount would mean to a resident of the foreign country, that is, if the money is consumed there, not here. Both questions—the purchasing power of some amount of money in your home country, and the purchasing power abroad—provide the same answer if prices abroad and at home are on average the same once they have been converted into the same currency. This situation is known as (Absolute) Purchasing Power parity (APPP). As we will illustrate below, APPP does not hold in reality, with deviations becoming more important the more different the two countries are in terms of location or economic development.

3.4.1 The PPP rate

To have a more reliable feel for what a given amount of foreign money really means locally, one needs for each country a number called the price level, which we denote by Π (at home, and in HC), and Π^* (abroad, and in FC), respectively. A price level is an absolute amount of currency—not an index number—needed to buy a standard consumption bundle. Computing price levels for different countries makes sense only if the consumption bundle whose cost is being measured is the same across countries. In a simple economy where fast food would be the only commodity, the bundle may be one soda, one burger, one fries (medium), a salad and a coffee—let's call this the BigMeal. We simply jot down the prices of the components abroad and at home, and tot them up into price levels for BigMeals.

Any differences in price levels, after conversion into a common currency, would make a simple conversion of a FC amount into the HC rather misleading if translated price levels are very diffferent:

Example 3.14

You often chat with a friend living in the Republic of Freedonia where, since the presidency of Groucho Marx, the currency is the Freedonian Crown (FDK). Let S_t = USD/FDK 0.010. You earn USD 50 per unit of time, your Freedonian friend 2000 FDK. What does that income really mean if the standard consumption bundle, our BigMeal, costs USD 5 here and FDK 250 in Freedonia?

- At the spot rate of USD/FDK 0.010, your friend seems to earn only $2000 \times 0,010$ = USD 20, suggesting that she is 60 percent worse off than you.

- But this ignores price differences. What you "really" earn is 50/5 = 10 BigMeals, while your friend makes 2000/250 = 8 BigMeals. That is, your friend is "really" almost as well off as you are.

What's the PPP rate?

To buy 8 BigMeals at home, you would need $8 \times 5 =$ USD 40. So one way to summarize the situation is that FDK 2000 means as much to your friend abroad as USD 40 means here, to you. The USD 40 is called the translation of FDK 2000 using the Purchasing Power Parity rate rather than the nominal rate, and the implied PPP rate is the $40/2000 = 0.020$ USD per FDK, the ratio of the two price levels.

Let's generalize. Suppose you want to have a feel for what a FC amount Y^* "really" means to a foreigner. The question can be made more precise as follows: Give me a HC amount $\hat{Y}$ such that its purchasing power here, $\hat{Y}/\Pi_t$, equals the purchasing power abroad of the original amount, Y^*/Π_t^*:

$$
\begin{aligned}
\text{Find } \hat{Y} \text{ such that } \frac{\hat{Y}}{\Pi_t} &= \frac{Y^*}{\Pi_t^*} \\
\Rightarrow \hat{Y} &= \frac{\Pi_t}{\Pi_t^*} Y^* && (3.9) \\
&= \hat{S}_t^{\text{PPP}} Y^*, \text{ where } \hat{S}_t^{\text{PPP}} \stackrel{\text{def}}{=} \frac{\Pi_t}{\Pi_t^*} && (3.10)
\end{aligned}
$$

So we can always compute the PPP rate as the ratio of the two price levels. For example, your friend's foreign amount (FDK 2000) could have been translated at the PPP rate, $5/250 = 0.020$, which would have told you immediately that her income buys as much (in Freedonia) as USD 40 buys here.

Example 3.15

End 2006, the CIA Factbook (http://www.cia.gov/cia/publications/factbook/geos/rs.html) assessed Russia's 2005 GDP at 1.589 USD trillion using the PPP rate, and at 740.7 billion using the nominal official rate. What is the explanation: are prices lower in Russia than in the US, or is it the inverse?

For China, the then figures were (purchasing power parity:) USD 8.883 trillion and (official exchange rate:) USD 2.225 trillion), a ratio of about four to one instead of Russia's two to one. Which country, then, has the lower price level?

DoItYourself problem 3.2

Check that the PPP rate has dimension HC/FC.

The IMF and the World Bank, for instance, often use PPP rates rather than the regular ("nominal") rate to translate foreign GDP's or incomes or government

Table 3.4: **PPP rates based in Big-Mac prices from *The Economist***

	currency	local price	actual value of $	PPP rate of $	real rate of $	actual value in $	PPP rate in $	real rate in $
China	yuan	10.5	8.03	3.39	2.371	0.125	0.295	0.422
Macau	pacata	11.1	7.99	3.58	2.231	0.125	0.279	0.448
Malaysia	ringgit	5.5	3.63	1.77	2.046	0.275	0.564	0.489
Hong Kong	dollar	12	7.75	3.87	2.002	0.129	0.258	0.499
Indonesia	rupia	14600	9325	4709.68	1.980	0.000	0.000	0.505
Philippines	peso	85	52.6	27.42	1.918	0.019	0.036	0.521
Paraguay	guarani	9000	5505	2903.23	1.896	0.000	0.000	0.527
Egypt	pound	9.5	5.77	3.06	1.883	0.173	0.326	0.531
Ukraine	hryvna	8.5	5.05	2.74	1.842	0.198	0.365	0.543
Moldava	leu	23	13.2	7.42	1.779	0.076	0.135	0.562
Uruguay	peso	42.3	23.9	13.65	1.752	0.042	0.073	0.571
Russia	ruble	48	27.1	15.48	1.750	0.037	0.065	0.571
Dominican Rep	peso	60	32.6	19.35	1.684	0.031	0.052	0.594
Sri Lanka	rupee	190	103	61.29	1.681	0.010	0.016	0.595
Honduras	lempira	35.95	18.9	11.60	1.630	0.053	0.086	0.614
Bulgaria	lev	2.99	1.54	0.96	1.597	0.649	1.037	0.626
Slovakia	koruna	58	29.5	18.71	1.577	0.034	0.053	0.634
Poland	zloty	6.5	3.1	2.10	1.478	0.323	0.477	0.676
Thailand	baht	60	28.4	19.35	1.467	0.035	0.052	0.682
South Africa	rand	13.95	6.6	4.50	1.467	0.152	0.222	0.682
Pakistan	rupee	130	60.1	41.94	1.433	0.017	0.024	0.698
Venezuela	bolivar	5701	2630	1839.03	1.430	0.000	0.001	0.699
Costa Rica	colon	1130	510	364.52	1.399	0.002	0.003	0.715
Japan	yen	250	112	80.65	1.389	0.009	0.012	0.720
Singapore	dollar	3.6	1.59	1.16	1.369	0.629	0.861	0.730
Guatemala	quetzal	17.25	7.59	5.56	1.364	0.132	0.180	0.733
Argentina	peso	7	3.06	2.26	1.355	0.327	0.443	0.738
Georgia	lari	4.15	1.8	1.34	1.345	0.556	0.747	0.744
Taiwan	dollar	75	32.1	24.19	1.327	0.031	0.041	0.754

	currency	local price	actual value of $	PPP rate of $	real rate of $	actual value in $	PPP rate in $	real rate in $
Estonia	kroon	29.5	12.3	9.52	1.293	0.081	0.105	0.774
Saudi Arabia	riyal	9	3.75	2.90	1.292	0.267	0.344	0.774
Lithuania	litas	6.5	2.69	2.10	1.283	0.372	0.477	0.779
Australia	dollar	3.25	1.33	1.05	1.269	0.752	0.954	0.788
UAE	dirham	9	3.67	2.90	1.264	0.272	0.344	0.791
Latvia	lats	1.35	0.55	0.44	1.263	1.818	2.296	0.792
Mexico	peso	29	11.3	9.35	1.208	0.088	0.107	0.828
Colombia	peso	6500	2504	2096.77	1.194	0.000	0.000	0.837
Croatia	kuna	15	5.72	4.84	1.182	0.175	0.207	0.846
South Korea	won	2500	952	806.45	1.180	0.001	0.001	0.847
Czech Rep	koruna	59.05	22.1	19.05	1.160	0.045	0.052	0.862
Fiju	dollar	4.65	1.73	1.50	1.153	0.578	0.667	0.867
Hungary	forint	560	206	180.65	1.140	0.005	0.006	0.877
Turkey	lire	4.2	1.54	1.35	1.137	0.649	0.738	0.880
New Zealand	dollar	4.45	1.62	1.44	1.129	0.617	0.697	0.886
Slovenia	tolar	520	189	167.74	1.127	0.005	0.006	0.888
Aruba	florin	4.95	1.79	1.60	1.121	0.559	0.626	0.892
Brazil	real	6.4	2.3	2.06	1.114	0.435	0.484	0.898
Morocco	dirham	24.5	8.71	7.90	1.102	0.115	0.127	0.907
Peru	new sol	9.5	3.26	3.06	1.064	0.307	0.326	0.940
Chile	peso	1560	530	503.23	1.053	0.002	0.002	0.949
United states	dollar	3.1	1	1.00	1.000	1.000	1.000	1.000
Canada	dollar	3.52	1.12	1.14	0.986	0.893	0.881	1.014
Britain	pound	1.94	0.532	0.63	0.850	1.880	1.598	1.176
Euroland	euro	2.94	0.781	0.95	0.824	1.280	1.054	1.214
Sweden	krona	33	7.28	10.65	0.684	0.137	0.094	1.462
Denmark	krone	27.75	5.82	8.95	0.650	0.172	0.112	1.538
Switzerland	franc	6.3	1.21	2.03	0.595	0.826	0.492	1.680
Iceland	kronur	459	72	148.06	0.486	0.014	0.007	2.056
Norway	kroner	43	6.1	13.87	0.440	0.164	0.072	2.274

Source Based on data from *The Economist*, May 26, 2006

budgets. Also newspapers or magazines start adopting this approach. Lastly, the PPP rate also serves as a benchmark for the nominal rate. Many economists feel that, in the long run, the nominal rate for two similar economies should loosely fluctuate around the PPP rate, and never wander off very far above or below it. Let's see how well this theory fares, empirically.

PPP in reality

In Table 3.4 we take *The Economist*'s favorite consumption bundle, the Big Mac,[17] and we compute PPP rates for 59 countries—once in USD (a New Yorker should get 0.295 dollars to be as happy as a Beijinger with one extra yuan) and once in non-US currency (a Beijinger should get 3.39 yuan to be as happy as a New Yorker with with one extra dollar). You see that countries where the Mac has a high local price have, of course, low PPP rates but also tend to have low actual exchange rates. Figure 3.14 shows this graphically. To "shrink" the outliers and give the smaller numbers more space we plot the log of the actual against the log of the PPP rate.

[17] Based on data from *The Economist*, May 26, 2006

Figure 3.14: **Log PPP *v* log actual rates, HC/USD**

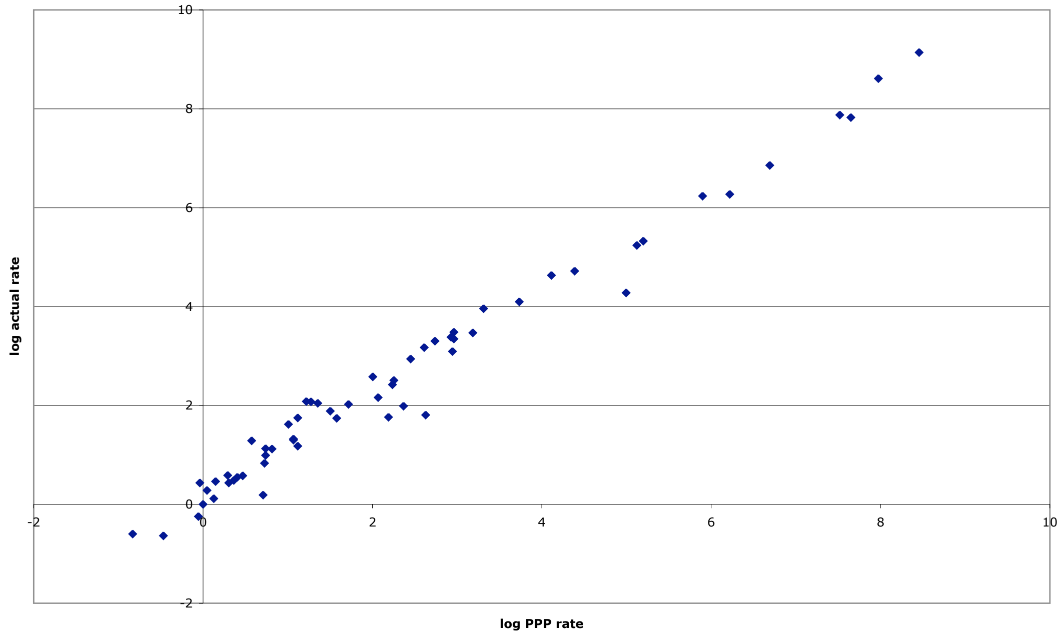

Source Based on data from *The Economist*, May 26, 2006

(This explains why there are negative rates: numbers below unit produce negative logs). There obviously is a very strong link.

DoItYourself problem 3.3

Knowing that the BigMac costs 3.10 in the US and 155 in Freedonia, and that the spot rate is 100 Crowns per dollar, complete Freedonia's PPP rates in the table:

	curren-cy	local price	actual value of $	PPP rate of $	real rate of $	actual value in $	PPP rate in $	real rate in $
Freedonia	korona	155	100					

.

But a closer look at the table reveals big relative deviations, which are hard to spot from a log graph dominated by outliers. Kicking out the 20 highest cases so as to be able to forego logs, this time, we get Figure 3.15. Note how the observations tend to be above the equality line (where actual = PPP): the dollar tends to be too expensive, by BigMac PPP standards. Yet there are also important deviations below the 45-degree line, where the slope of the ray through the dot is even below 0.5 in one case. The slope of this ray is called the *real exchange rate*, to which we now turn.

Figure 3.15: **PPP *v* actual rates, HC/USD**

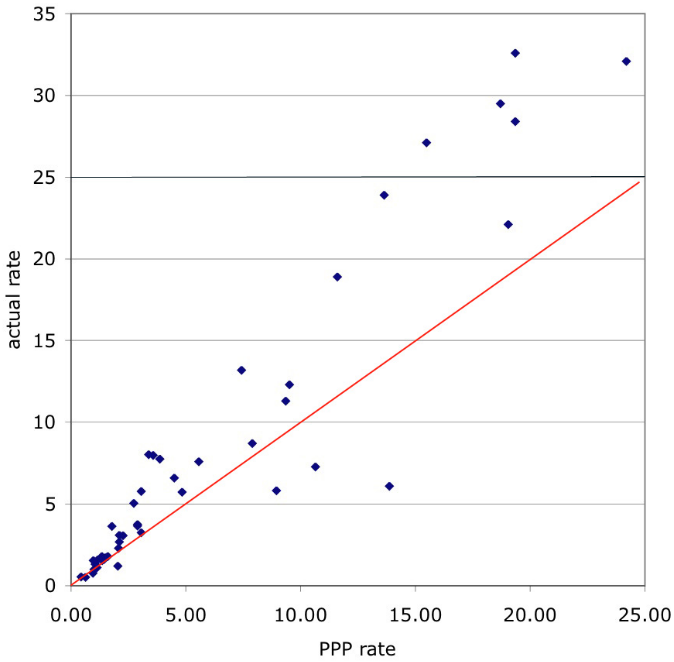

Source Based on data from *The Economist*, May 26, 2006

3.4.2 Commodity Price Parity

A concept used in textbooks is Commodity Price Parity (CPP). It is said to hold when translated prices for an individual good are equalized across two countries:

$$\text{CPP holds if } P_{j,t} = S_t \times P^*_{j,t}, \tag{3.11}$$

with j referring to an individual good, and P_j (P^*_j) referring to its price at home, in HC (abroad, in FC). In fact, all the Big Mac evidence shown thus far is about CPP rather than PPP, a distinction that *The Economist* tends to gloss over.

CPP would hold if trading were costless and instantaneous. Obviously, in reality it does not work across the board; for commodities it is not too bad an approximation (within the bounds created by transportation costs and the like), but for consumer goods it is essentially a joke.

PPP in the true sense—*i.e.* for a bundle of goods—would clearly hold if CPP held for every individual good, or if deviations from CPP washed out after averaging across many goods. As we have seen, this is not really the case; apparently, too many deviations from CPP turn out to be in the same direction, suggesting there is a common force behind them. Forget CPP.

3.4.3 The Real Exchange Rate and (Deviations from) Absolute PPP

The real exchange rate (RER) is a measure of how far the nominal rate differs from the PPP one: it simply is the nominal exchange rate divided by the PPP counterpart.

Example 3.16

In our Freedonian story, the nominal rate was 0.010 USD/FDK while the PPP rate was 0.020 USD/FDK; thus, the real rate was 0.5—a large deviation from unity, but not uncommon between two very different economies.

The real rate is a dimensionless number—[HC/FC] divided by [HC/FC]. In a way, it simply tells us what the ratio is of the *translated* price levels:

$$RER_t \stackrel{\text{def}}{=} \frac{S_t}{\hat{S}_t^{\text{PPP}}} \tag{3.12}$$

$$= \frac{S_t \times \Pi_t^*}{\Pi_t}, \text{ from (3.10).} \tag{3.13}$$

Again, in the example one can find the RER for the FDK against the USD by translating into USD the foreign price of the BigMeal, FDK $250 \times 0.010 =$ USD 2.5, and divide it by the domestic price level, 5, which gets us $2.5/5 = 0.5$. Thus, the RER rate tells you how much cheaper (if RER < 1) or more expensive (if RER > 1) the foreign country is. A country with a below-unity real rate would be a nice place to spend your domestic income, or could be an attractive base to export from, but may not be the best place to export to. These are very different questions than the one answered by the PPP rate.

Obviously, if the real rate equals unity, both countries have the same price level. If that is true, Absolute PPP is said to hold:

$$\text{Absolute PPP holds if } RER_t = 1 \Leftrightarrow S_t = S_t^{\text{PPP}} \Leftrightarrow S_t \times \Pi_t^* = \Pi_t \tag{3.14}$$

In Figure 3.16, and in Table 3.4 the countries have been ranked on the basis of the real rate. Two observations stand out. First, there is a five-to-one ratio between the most and least expensive countries, Norway and China. So deviations from PPP are big. Second, there is a system to it, to some extent: undervalued currencies tend to be developing ones, and overvalued ones developed. (The fact that thus USD is not top is anomalous, in this view. The long-lasting deficit in its CA may be one reason). The (imperfect but strong) relation between real rate and degree of economic development is discussed in Chapter 10.

DoItYourself problem 3.4

Norway is most expensive. Identify the dot, in Figure 3.15, that corresponds to Norway.

Figure 3.16: **Real rates based on Big-Mac prices from *The Economist***

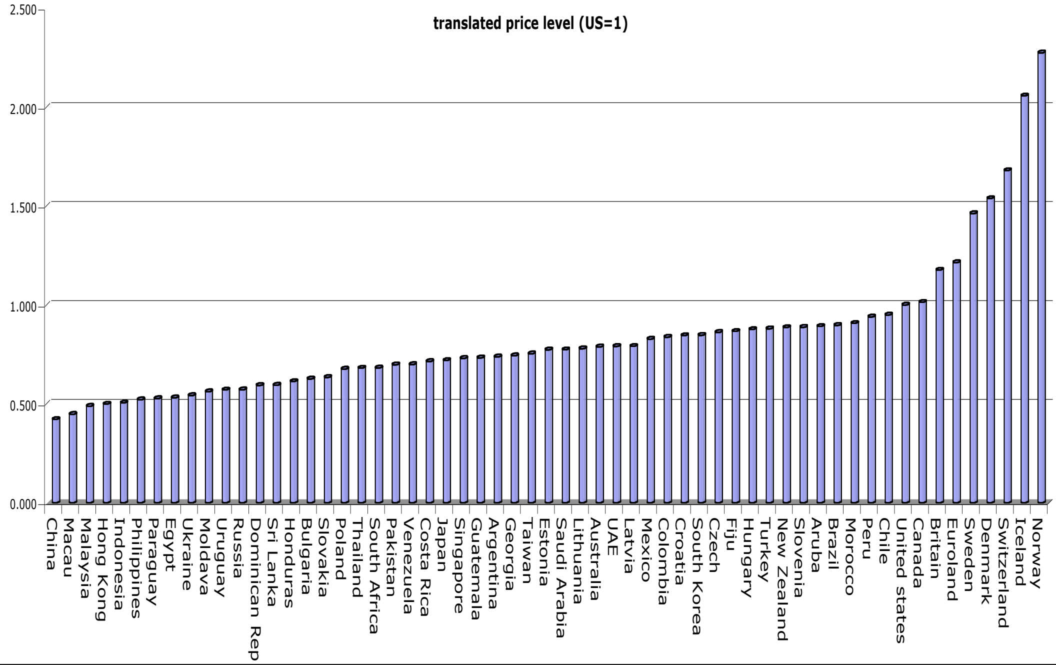

Source Based on data from *The Economist*, May 26, 2006

3.4.4 The Change in the Real Rate and (deviations from) Relative PPP

For most of the time since the 80s, Japan has had a real rate above unity: it was a more expensive place to spend a dollar than the US or Europe. Sometimes one would be interested in whether the country's situation has worsened or improved. That is, has the real rate increased or decreased (as distinct from the issue of whether its level is above unity or not)?

To measure this, one can simply compute the RER's percentage change. Not surprisingly, the percentage change in the RER is determined by the percentage changes in the spot rate and the price levels—the inflation rates:

Example 3.17

Q. Suppose that 5 years ago the FDK traded at USD/FDK 0.012, and the price levels were USD 4 in the US and FDK 250 abroad. (So, with current price levels being 5 and 250, respectively, inflation was 25 percent in the US, and zero in Freedonia.) Recalling that the current RER is $250 \times 0.010/5 = 0.5$, how much did the RER change since then?

A. The old RER was $250 \times 0.012/4 = 0.75$; so the rate changed by $(0.50-0.75)/0.75 = -0.33$, that is, minus 33 percent. There was real depreciation of the Crown—that is, Freedonia became cheaper over time—because the FDK went down *and* because inflation in Freedonia was lower than US.

Below, we first show the general relation between the percentage change in the RER and the changes in the nominal rate, and then a first-order approximation that is occasionally used:

$$\begin{aligned} \text{percentage change in the RER} &= (1+s_{t_0,t})\frac{1+\mathit{infl}^*_{t_0,t}}{1+\mathit{infl}_{t_0,t}} - 1 && (3.15)\\ &\approx s_{t_0,t} + [\mathit{infl}^*_{t_0,t} - \mathit{infl}_{t_0,t}] && (3.16)\end{aligned}$$

where $s_{t_0,t}$ is the simple percentage change in the spot rate S between times t_0 and t while $\mathit{infl}_{t_0,t}$ and $\mathit{infl}^*_{t_0,t}$ denote inflation at home and abroad, respectively, over the same time window. The first-order approximation works well if both inflation rates are low. This is not the case in our Freedonian example:

Example 3.18

In our above story, foreign inflation was zero, US inflation 25 percent, and the exchange rate changed by minus one-sixth; so the RER changes by

$$(1-1/6)\frac{1+0.00}{1+0.25} - 1 = 0.66667 - 1 = -1/3,$$

as computed directly before. In contrast, the first-order approximation would have predicted a change of –1/6 – 0.25 = –41.67 percent rather than –33.33 percent. The error is nontrivial because in this example the exchange-rate change and one inflation rate, the US one, are far from zero.

If the RER is constant—whatever the level— then Relative PPP (RPPP) is said to hold; and the percentage change in the RER is a standard measure of deviations from RPPP. An RPPP deviation is most often resorted to if the RER itself cannot be computed because price-level data are missing. If, indeed, absolute price levels for identical bundles are not available, there is no way to compute which of the two countries is the cheaper one. But one can still have an idea whether the RER went up or down if one estimates the inflation rates from the standard Consumption Price Indices (CPI's) rather than the price levels. A CPI is a relative number vis-a-vis a base period, and the consumption bundle is typically tailored to the country's own consumption pattern rather than being a common, internationally representative bundle of goods. Still, in most cases this makes little difference to the inflation rates.

The RPPP rate relative to some chosen base period t_0 is the level of the current

Figure 3.17: **[Actual Rate]/[RPPP Rate] against the usd, 1965=1.00**

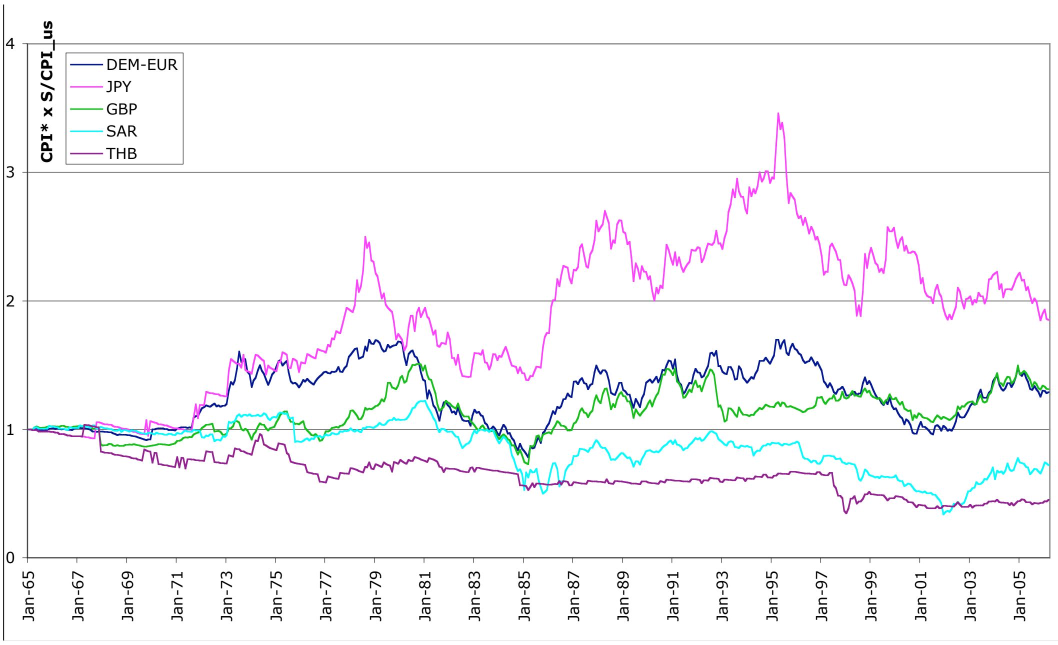

Source Underlying data are from Datastream

rate that keeps the RER at the same level as in the base period:

$$[\text{RPPP rate vis-a-vis } t_0] = S_t^{RPPP,t_0} = S_{t_0}\frac{1+infl_{t_0,t}}{1+infl^*_{t_0,t}}; \tag{3.17}$$

$$\text{Relative Real Rate vis-a-vis } t_0 = \frac{S_t}{S_t^{RPPP,t_0}},$$
$$= \frac{S_t}{S_{t_0}}\frac{1+infl^*_{t_0,t}}{1+infl_{t_0,t}}, \tag{3.18}$$

which is unity plus the change in the real rate except that we use each country's CPI inflation (or some similar index) rather than the change in the absolute price of an internationally common basket. In pre-EUR days, the EC or EU ministers of EMR[18] countries used the RPPP norm when devaluations were negotiated. They went back to the time of the last re-alignment, and corrected that base-period level for the accumulated inflation differential since then, as in Equation [3.17]. But the main use of the RPPP for business is that it tells us whether a country has become cheaper, or more expensive, relative to another one. Cheapening countries are good if they are your production centers or your favorite holiday resort, but bad if they are the

[18]Exchange Rate Mechanism—the arrangement that kept members' crossrates stable. See Chapter 2.

Figure 3.18: **RPPP v actual rates against USD, 1965=1.00**

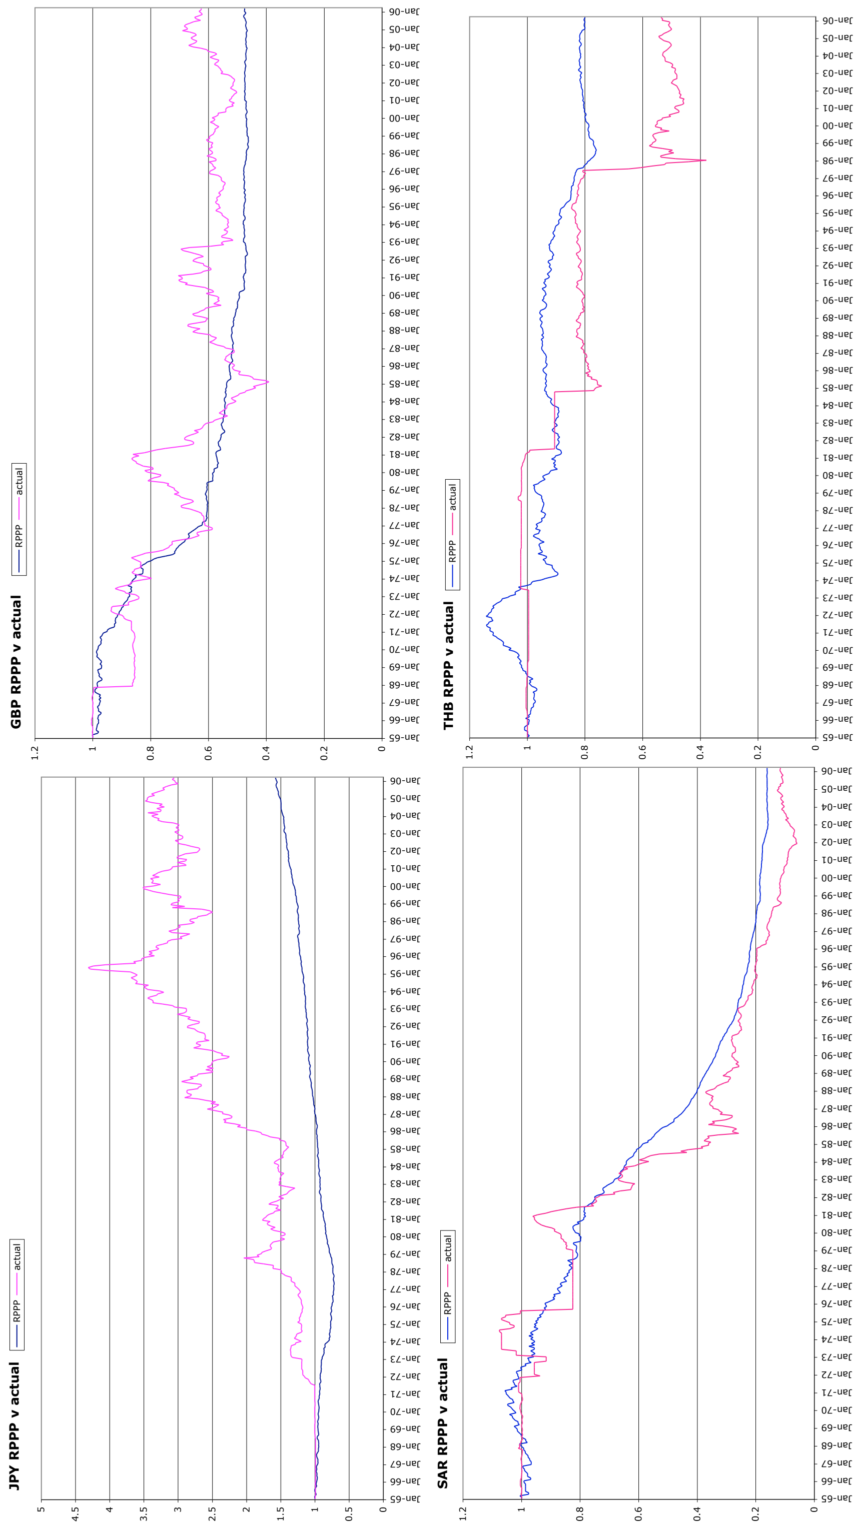

Source Underlying data are from Datastream markets where you sell your output.

For this reason, deviations from RPPP are important. Are they large? Figure 3.17 shows time-series data, taking Jan 1965 as the base period, on relative real rates against USD, for the DEM-EUR, JPY, GBP, SAR, and THB. We note four facts.

- First, there are huge swings in the medium run, with the real rate appreciating by 50% and then going back—and occasionally even doubling or halving— in a matter of years not decades. Imagine you being caught in this as an exporter.

- Second, in the short run there is lots of inertia: once the rate is above its mean, it tends to stay there for years. Statistical analysis shows that the average half life is three to five years, meaning that it takes three to five years, on average, for a deviation to shrink to half its original size. Thus, when you get into a bad patch, you can expect that this will be a matter of years rather than weeks or months.

- Third, when we look at the RPPP rates and the actual ones separately (Figure 3.18), we see that, almost always, in the short run most of the variation in the real rate stems from the nominal rate; the RPPP rate is usually smooth relative to the actual, except of course under a fixed-exchange-rate regime (see graphs) and in hyperinflation cases (not shown). The fall, rise, & fall of the USD against the DEM under presidents Carter and Reagan had nothing to do with inflation. In a way, that's good, because there are good hedge instruments against swings in nominal rates. Hedging nominal rates, in the short run, almost stabilizes the real rate too.

- Even though deviations between actual and RPPP rates are huge, there often does seem to be a link, in the long run. As a result, the long-run variability of the inflation-corrected rate is somewhat lower than that in the nominal rate.

- A last fact, impossible to infer from the graphs but to be substantiated in Chapter 10, is that changes in both nominal and real exchange rates seem hard to predict.

Should you care? If exchange risk would just lead to capitals gains and losses on assets or liabilities denominated in FC, most (but not all) firms would be able to shrug it off as a nuisance, perhaps, but no more than that. However, there is more: real-rate moves may also make your production sites incompetitive or your export markets unprofitable, and it is harder for a firm to just shrug this off. Another implication worth mentioning is that when two investors from different countries hold the same asset, they will nevertheless realize different real returns if the real exchange rate is changing—which it does all the time. Thus, exchange risk undermines one of the basic assumptions of the CAPM, namely that investors all agree on expected returns and risks. These implications explain why exchange risk gets so much attention in this text.

3.5 CFO's Summary

In this chapter, we have seen how spot markets work. From the treasurer's point of view, one immediately interesting aspect is the possibility for arbitrage and shopping around.

- Arbitrage consists of buying and immediately reselling (or vice versa), thus taking no risk and engaging no capital. One could try to do this across market makers (for one particular exchange rate) or in a triangular way. In practice, the likelihood of corporate treasurers finding such a riskless profit opportunity is tiny. Arbitrage by traders in the wholesale market eliminates this possibility almost as quickly as it arises. In addition, most firms deal in the retail market, where spreads are relatively wide.

- Shopping around consists of finding the best route for a particular transaction. In contrast to arbitrage, shopping around may work—not in the sense of creating large profits, but in the sense of saving on commissions or getting marginally better rates. It is generally worth calling a few banks for the best rate when you need to make a large transaction. And it may pay to compute a triangular cross rate, especially through routes that involve heavily traded currencies like the USD or the EUR. Doing such a computation could enable corporate treasurers to find cheaper routes for undertaking transactions as compared to direct routes.

The spot rate is, by definition, the right number to use if you need to do an actual transaction. But for other purposes, other exchange-rate concepts are quite useful:

- The PPP rate is the ratio of the two price levels. Translating foreign income numbers or investment budgets at this rate tells you what the foreign figures really mean to locals—but expressed in terms that are familiar to you.

- The *real exchange rate* (or the deviation from Absolute PPP) is the ratio of the *translated* price levels. It tells you which country is more expensive. This is relevant if you want to evaluate a country as a destination for exports, or a source of imports, or a place to live or produce.

- Both the above concepts require data on price levels, which are not available for all countries. Often one makes do with the *deviation from RPPP* relative to a given base period, which estimates to what extent the real exchange rate has changed since then.

There is a clear, but imperfect relation between actual rates and PPP rates: countries that have gone through high-inflation episodes and, thus, ended up with high nominal prices for all goods, pay high nominal prices for currencies too. But the relation is far from one-to-one: real rates can be five to one (Norway against China, for instance, in the BigMac data set). Also, there is a lot of variability over time, making countries more attractive or unattractive as production centers or markets. Most of that variability comes from the nominal exchange rate: inflation contributes little, except for hyper-inflation episodes (with inflation rates measured in 100s or 1000s per month). Thus, currency risk affects contractual cash flows fixed in FC but also the operations of a firm. It even messes up the CAPM because real exchange risk means that investors from different countries no longer perceive asset returns and risks in the same way.

What are the implications for the CFO? You should remember, first, that variations in the real exchange rate are long-memory events and can be vast. So they can have a big impact on how and where you should produce, and may even force you to change your fundamental strategy. All this comes on top of a shorter-run effect, of course: variations in exchange rates cause capital gains and losses on FC-denominated contractual claims and liabilities.

Your instinctive reaction may be that the firm should try to reduce the impact of these changes. This may be too fast, though: we first need to determine whether any such "hedging" policy really adds value. To be able to answer this question, we need to understand how the hedge instruments work: forwards, futures, swaps, and options. A good knowledge of these derivatives is, of course, also required to make an informed choice among the available hedge instruments. This is what the remainder of Part II is about. We begin with forward markets.

References

Grabbe, O., 1995. International Financial Markets, 3rd Edition, New York: Elsevier

Deardorff , a. V., 1979. One-Way Arbitrage and Its Implications for the Foreign Exchange Markets. The Journal of Political Economy, Vol. 87 Issue 2, p351-65

3.6 TekNotes

Technical Note 3.1 What's wrong with the FC/HC convention, in a textbook?

In the text just below Example 3.2 we claimed that using the FC/HC convention would mean all the familiar formulas from Finance would have to be abandoned. Here's this message in math. Let r^* denote the risk-free interest rate earned on FC, and $\tilde{S}_1$ the (random) future value, in HC of one unit of FC. If you buy one unit of FC, you'll have $1 + r^*$ of them next period, worth $\tilde{S}_1(1 + r^*)$ in HC. Standard finance theory then says that the current price, S_0, should be the future value discounted at a rate $\mathrm{E}(\tilde{r}_S)$ that takes into account this risk of $\tilde{S}_1$:

$$S_0 = \frac{\mathrm{E}(\tilde{S}_1)(1 + r^*)}{1 + \mathrm{E}(\tilde{r}_S)}. \tag{3.19}$$

This looks quite normal and well behaved. Now look at what would happen if we had used the inverse rate, $X := S^{-1}$, and if we wanted a theory about how X_0 is set. First substitute $X = S^{-1}$ into the equation and then solve for X_0:

$$1/X_0 = \frac{\mathrm{E}(1/\tilde{X}_1)(1 + r^*)}{1 + \mathrm{E}(\tilde{r}_S)} \Rightarrow X_0 = \frac{1 + \mathrm{E}(\tilde{r}_S)}{\mathrm{E}(1/\tilde{X}_1)(1 + r^*)}.$$

All connection with finance is gone. The discount rate is on top (?!), and the expectation is below, and the expectation is about the inverse of X. Clearly, this makes no sense in a finance textbook.

3.7 Test Your Understanding

3.7.1 Quiz Questions

1. Using the following vocabulary, complete the text: forward; market maker or broker; shopping around; spot; arbitrage; retail; wholesale.

   "When trading on the foreign exchange markets, the Bank of Brownsville deals with a (a) on the (b) tier, while an individual uses the (c) tier. If the bank must immediately deliver EUR 2 million to a customer, it purchases them on the (d) market. However, if a customer needs the EUR in three months, the bank buys them on the (e) market. In order to purchase the EUR as cheaply as possible, the bank will look at all quotes it is offered to see if there is an opportunity for (f). If the bank finds that the quotes of two market makers are completely incompatible, it can also make a risk-free profit using (g)."

2. From a Canadian's point of view, which of each pair of quotes is the direct quote? Which is the indirect quote?

   (a) CAD/GBP 2.31; GBP/CAD 0.43

   (b) USD/CAD 0.84; CAD/USD 1.18

   (c) CAD/EUR 1.54; EUR/CAD 0.65

3. You are given the following spot quote: EUR/GBP 1.5015-1.5040

   (a) The above quote is for which currency?

   (b) What is the bid price for EUR in terms of GBP?

4. You read in your newspaper that yesterday's spot quote was CAD/GBP 2.3134-2.3180.

   (a) This is a quote for which currency?

   (b) What is the ask rate for CAD?

   (c) What is the bid rate for GBP?

5. A bank quotes the following rates. Compute the EUR/JPY bid cross-rate (that is, the bank's rate for buying JPY).

	Bid	Ask
EUR/CAD	0.64	0.645
CAD/JPY	0.01	0.012

6. A bank quotes the following rates: CHF/USD 2.5110-2.5140 and JPY/USD 245–246. What is the minimum JPY/CHF bid and the maximum ask cross rate that the bank would quote?

7. A bank is currently quoting the spot rates of EUR/USD 1.3043-1.3053 and NOK/USD 6.15-6.30. What is the lower bound on the bank's bid rate for the NOK in terms of EUR?

8. Suppose that an umbrella costs USD 20 in Atlanta, and the USD/CAD exchange is 0.84. How many CAD do you need to buy the umbrella in Atlanta?

9. Given the bid-ask quotes for JPY/GBP 220-240, at what rate will:

   (a) Mr. Smith purchase GBP?
   (b) Mr. Brown sell GBP?
   (c) Mrs. Green purchase JPY?
   (d) Mrs. Jones sell JPY?

True or false? Indicate the correct statement(s).

1. CPP says that you can make a risk-free profit by buying and selling goods across countries.

2. CPP implies causality. It states that foreign prices are determined by domestic prices and other factors such as production costs, competitive conditions, money supplies, and inflation rates.

3. In order for a firm not to be affected by real exchange risk, CPP must hold not only for the goods a firm produces but also for all production inputs, and for the prices of complementary and substitute goods.

4. The equilibrium exchange rate suggested by the Absolute Purchasing Power Parity hypothesis depends on the relative relationship between the prices of a representative consumption bundle in the currencies of two countries.

5. Your purchasing power is the number of representative consumption bundles that you can buy.

6. The real effective exchange rate is the price of an average foreign consumption bundle in units of domestic currency.

7. Relative PPP shows how a consumer's purchasing power changes over time.

8. Absolute PPP may hold even when Relative PPP does not because absolute PPP looks at levels at a specific point in time, and levels are always comparable regardless of the composition of the consumption bundle.

9. Given the empirical evidence on the correlation between the nominal and real exchange rate, it is possible to use the nominal financial instruments to hedge real exchange risk.

10. Purchasing Power Parity is based on the idea that the demand for a country's currency is derived from the demand for that country's goods as well as the currency itself.

Multiple-Choice Questions Choose the correct answer(s).

1. CPP may not hold because:
   (a) the prices for individual goods are sticky.
   (b) transaction costs increase the bounds on deviations from CPP, making it more difficult to arbitrage away price differences.
   (c) quotas and voluntary export restraints limit the ability to arbitrage across goods markets.
   (d) parallel imports lead to two different prices for the same good.
   (e) the prices of tradable goods fluctuate too much, which makes it difficult to take advantage of arbitrage opportunities.

2. Absolute Purchasing Power Parity may not hold when:
   (a) the prices of individual goods in the consumption bundle consistently deviate from CPP across two countries.
   (b) the consumption bundles of different countries are not the same.
   (c) the prices for individual goods are sticky.
   (d) there are tariffs, quotas, and voluntary export restraints.
   (e) competition is perfect.

3. Relative Purchasing Power Parity is relevant because:
   (a) empirical tests have shown that Absolute PPP is always violated, while Relative PPP is a good predictor of short-term exchange rate exposure.
   (b) consumption bundles are not always comparable across countries.
   (c) price levels are not stationary over time.
   (d) investors care about the real return on their international portfolio investments.
   (e) investors care about the nominal return on their international portfolio investments.

3.7.2 Applications

1. You have just graduated from the University of Florida and are leaving on a whirlwind tour to see some friends. You wish to spend USD 1,000 each in Germany, New Zealand, and Great Britain (USD 3,000 in total). Your bank offers you the following bid-ask quotes: USD/EUR 1.304-1.305, USD/NZD 0.67-0.69, and USD/GBP 1.90-1.95.

(a) If you accept these quotes, how many EUR, NZD, and GBP do you have at departure?

(b) If you return with EUR 300, NZD 1,000, and GBP 75, and the exchange rates are unchanged, how many USD do you have?

(c) Suppose that instead of selling your remaining EUR 300 once you return home, you want to sell them in Great Britain. At the train station, you are offered GBP/EUR 0.66-0.68, while a bank three blocks from the station offers GBP/EUR 0.665-0.675. At what rate are you willing to sell your EUR 300? How many GBP will you receive?

2. Abitibi Bank quotes JPY/EUR 155-165, and Bathurst Bank quotes EUR/JPY 0.0059-0.0063.

   (a) Are these quotes identical?

   (b) If not, is there a possibility for shopping around or arbitrage?

   (c) If there is an arbitrage opportunity, how would you profit from it?

The following spot rates against the GBP are taken from the Financial Times of Friday, February 2, 2007. Use the quotes to answer the questions in Exercises 3 through 5.

Country	Code	midpoint	change	spread
Czech Rep	CZJ	42.7945	+0.1868	616–273
Denmark	DKK	11.30929	+0.0289	065–119
Euro	EUR	1.5172	+0.0039	168–175
Norway	NOK	12.3321	+0.0394	263–379
Russia	RUB	52.1528	– 0.0368	376–679
Switzerland	CHF	2.4531	+0.0040	522–540
Turkey	YTL	2.7656	–0.0050	614–698

Note: Bid-ask spreads show only the last three decimal places. When the ask seems to be smaller then the bid, add 1000.

3. What are the bid-ask quotes for:

   (a) CZJ/GBP?

   (b) DKK/GBP?

   (c) EUR/GBP?

   (d) NOK/GBP?

4. What are the bid-ask quotes for:

   (a) GBP/CZJ?

   (b) GBP/DKK?

   (c) GBP/EUR?

   (d) GBP/NOK?

5. What are the cross bid-ask rates for:

   (a) RUB/CHF?
   (b) NOK/YTL?
   (c) DKK/EUR?
   (d) CZJ/CHF?

6. In Figure 2.8 I showed plots of the gold price and mentioned that, if we had corrected for inflation, then the 1980 price would be seen to be much above the current peak: obviously, the small percentage price rise of gold, between 1980 and 2007, must have been way below the percentage rise of the USCPI.

   (a) In the above we presumably use USCPI rate to deflate the USD prices. But is this result generalizable to all countries—is this conclusion necessarily also valid for Japanese or German investors? Why (not)?
   (b) If you think the result does not necessarily hold true elsewhere, what would you bet w.r.t. a hyper-inflator like Zimbabwe?: if inflation is much higher, then the real price of gold must have fallen even more—no?
   (c) What would guarantee identical real price paths in all countries: APPP, RPPP, or what?

Chapter 4

Understanding Forward Exchange Rates for Currency

In this chapter, we discuss forward contracts in perfect financial markets. Specifically, we assume that there are no transactions costs; there are no taxes, or at least they are non-discriminatory: there is but one overall income number, with all capital gains and interest earned being taxable and all capital losses and interest paid deductible; there is no default risk; and people act as price takers in free and open markets for currency and loans or deposits. Most of the implications of market imperfections will be discussed in later chapters; in this chapter we provide the fundamental insights that need to be mildly qualified later.

In Section 1, we describe the characteristics of a forward contract and how forward rates are quoted in the market. In Section 2, we show, with a simple diagram, the relationship between the money markets, spot markets, and forward markets. Using the mechanisms that enforce the Law of One Price, Section 3 then presents the Covered Interest Parity Theorem. Two ostensibly unconnected issues are dealt with in Section 4: how do we determine the market value of an outstanding forward contract, and how does the forward price relate to the expected future spot price. We wrap up in Section 5.

4.1 Introduction to Forward Contracts

Basics

Let us recall, from the first chapter, the definition of a forward contract. Like a spot transaction, a forward contract stipulates how many units of foreign currency are to be bought or sold and at what exchange rate. The difference with a spot deal, of course, is that delivery and payment for a forward contract take place in the future (for example, one month from now) rather than one or two working days from now, as in a spot contract. The rate that is used for all contracts initiated at time t and maturing at some future moment T is called the time-t forward rate for delivery date T. We denote it as $F_{t,T}$.

Like spot markets, forward markets are not organized exchanges, but over-the-counter (OTC) markets, where banks act as market makers or look for counterparts via electronic auction systems or brokers. The most active forward markets are the markets for 30 and 90 days, and contracts for 180, 270, and 360 days are also quite common. Bankers nowadays quote rates up to ten years forward, and occasionally even beyond that, but the very long-term markets are quite thin. Recall, lastly, that any multiple of thirty days means that, relative to a spot contract, one extra calendar month has to be added to the spot delivery date, and that the delivery date must be a working day. Thus, if day [$t + 2$ plus n months] is not a working day, we may move forward to the nearest working day, unless this would make us change months: then we'd move back.

Example 4.1

A 180-day contract signed on Thursday March 2, 2006 is normally settled on September 6. Why? The initiation day being a Thursday, the "spot" settlement date is Monday, March 6. Add 6 months; September 6, being a working day (Wednesday), then is the settlement date.

Market Conventions for Quoting Forward Rates

Forward exchange rates can be quoted in two ways. The most natural and simple quote is to give the actual rate, sometimes called the *outright* rate. This convention is used in, for instance, The *Wall Street Journal*, the *Frankfurter Allgemeine*, and the Canadian *Globe and Mail*. The *Globe and Mail* is one of the few newspapers also quoting long-term rates, as Table 4.1 shows.

In Table 4.1, the CAD/USD forward rate exceeds the spot rate for all maturities. Traders would say that the USD trades *at a premium*. Obviously, if the CAD/USD rate is at a premium, the USD/CAD forward rates must be below the USD/CAD spot rate; that is, the CAD must trade *at a discount*.

The second way of expressing a forward rate is to quote the difference between the outright forward rate and the spot rate—that is, quote the premium or discount. A forward rate quoted this way is called a swap rate.[1] Antwerp's *De Tijd*, or the London *Financial Times*, for example, used to follow this convention. Since both newspapers actually showed bid and ask quotes, we will postpone actual excerpts from these newspapers until the next chapter where spreads are taken into consid-

[1] Confusingly, the terms swap contract and swap rate can have other meanings, as we shall explain in Chapter 7.

Table 4.1: **Spot and Forward Quotes, Mid-market rates in Toronto at noon**

	(Outright)		(Swap rates) Premium or discount, in cents	
	CAD per USD	USD per CAD	CAD per USD	USD per CAD
U.S. Canada spot	1.3211	0.7569		
1monthforward	1.3218	0.7565	+0.07	–0.05
2 months forward	1.3224	0.7562	+0.13	–0.07
3 months forward	1.3229	0.7559	+0.18	–0.10
6 months forward	1.3246	0.7549	+0.35	–0.20
12 months forward	1.3266	0.7538	+0.55	–0.31
3 years forward	1.3316	0.7510	+1.05	–0.59
5 years forward	1.3579	0.7364	+3.68	–0.05
7 years forward	1.3921	0.7183	+7.10	–3.86
10 years forward	1.4546	0.6875	+13.36	–6.94

Source Globe and Mail .

eration. The rightmost two columns in Table 4.1 shows how *The Globe and Mail* quotes would have looked in swap-rate form. In that table, the sign of the swap rate is indicated by a plus sign or a minus sign. The *Financial Times* used to denote the sign as pm (premium) or dis (discount).

The origin of the term swap rate is the swap contract. In the context of the forward market, a swap contract is a spot contract immediately combined with a forward contract in the opposite direction.

Example 4.2

To invest in the US stock market for a few months, a Portuguese investor buys USD 100,000 at EUR/USD 1.10. In order to reduce the exchange risk, she immediately sells forward USD 100,000 for ninety days, at EUR/USD 1.101. The combined spot and forward contract—in opposite directions—is a swap contract. The swap rate, EUR/USD 0.1(cent), is the difference between the rate at which the investor buys and the rate at which she sells.

To emphasize the difference between a stand-alone forward contract and a swap contract, a stand-alone forward contract is sometimes called an outright contract. Thus, the two quoting conventions described above have their roots in the two types of contracts. Today, the outright rate and the swap rate are simply ways of quoting, used whether or not you combine the forward trade with a reverse spot trade.[2]

One key result of this chapter is that there is a one-to-one link between the swap

[2] Sometimes the swap rate is called the cost of the swap, but to financial economists that is a very dubious concept: at the moment the contract is initiated, both the spot and the forward part are zero-NPV deals, that is, their market value is zero. So the swap rate is not the cost of the swap in the same way a stock price measures the cost of a stock. It is more like an accounting concept of cost, in the style of the interest being the cost of a loan.

rate and the interest rates for the two currencies. To explain this relation, we first show how the spot market and the forward market are linked to each other by the money markets for each of the two currencies. But first we need to agree on a convention for denoting risk-free returns.

Our Convention for Expressing Risk-free Returns

We adopt the following terminology: the (effective) risk-free (rate of) return is the simple percentage difference between the initial, time-t value and the final, time-T value of a nominally risk-free asset over that holding period.

Example 4.3

Suppose that you deposit CLP 100,000 for four years and that the deposit will be worth CLP 121,000 at maturity. The four-year effective (rate of) return is:

$$r_{t,T} = \frac{121,000 - 100,000}{100,000} = 0.21 = 21 \text{ percent.} \tag{4.1}$$

You can also invest for nine months. Suppose that the value of this deposit after nine months is 104,200. Then the nine-month effective return is:

$$r_{t,T} = \frac{104,200 - 100,000}{100,000} = 0.042 = 4.2 \text{ percent.} \tag{4.2}$$

Of course, at any moment in time, the rate of return you can get depends on the time to maturity, which equals $T - t = 4$ years in the first example. Thus, as in the above examples, we always equip the rate of return, r, with two subscripts: $r_{t,T}$. In addition, we need to distinguish between the domestic and the foreign rate of return. We do this by denoting the domestic and the foreign return by $r_{t,T}$ and $r^*_{t,T}$, respectively.

It is important to understand that the above returns, 21 percent for four years and 4.2 percent for nine months, are not expressed on an annual basis. This is a deviation from actual practice: bankers always quote rates that are expressed on an annual basis. We shall call such a *per annum* ($p.a.$) percentage an *interest rate*. If the time to maturity of the investment or loan is less than one year, your banker will typically quote you a simple $p.a.$ interest rate. Given the simple $p.a.$ interest rate, you can then compute the effective return as:

$$r_{t,T} = [\text{time to maturity, in years}] \times [\text{simple } p.a. \text{ interest rate for that maturity}]. \tag{4.3}$$

Example 4.4

Suppose that the $p.a.$ simple interest rate for a three-month investment is 10 percent. The time to maturity, $T - t$, is 1/4 years. The effective return, then, is:

$$r_{t,T} = (1/4) \times 0.10 = 0.025. \tag{4.4}$$

The convention that we adopt in this text is to express all formulas in terms of effective returns, that is, simple percentage differences between end values and initial values. One alternative would be to express returns in terms of *per annum* simple interest rates—that is, we could have written, for instance, $(T-t)\,R_{t,T}$, (where capital R would be the simple interest on a *p.a.* basis) instead of $r_{t,T}$. Unfortunately, then all formulas would look more complicated. Worse, there are many other ways of quoting an interest rate in *p.a.* terms, such as interest with annual, or monthly, or weekly, or even daily compounding; or banker's discount; or continuously compounded interest. To keep from having to present each formula in many versions (depending on whether you start from a simple rate, or a compound rate, etc.), we assume that you have already done your homework and have computed the effective return from your *p.a.* interest rate. Appendix 4.6 shows how effective returns can be computed if the *p.a.* rate you start from is not a simple interest rate. That appendix also shows how returns should *not* be computed.

Thus, in this section, we will consider four related markets—the spot market, the forward market, and the home and foreign money markets. One crucial insight we want to convey is that any transaction in one of these markets can be replicated by a combination of transactions in the other three. Let us look at the details.

4.2 The Relation Between Exchange and Money Markets

We have already seen how, using the spot market, one type of currency can be transformed into another at time t. For instance, you pay home currency to a bank and you receive foreign currency. Think of one wad of HC bank notes being exchanged for another wad of FC notes. Or even better, since spot deals are settled second working days: think of a spot transaction as an exchange of two cheques that will clear two working days from now. As of now, we denote the amounts by HC and FC. To make clear that we mean amounts, not names, they are written as math symbols (full-sized and slanting), not as FC and FC, our notation for names of currency. Another notational difference between currency names and amounts is that FC and HC always get a time subscript. To emphasize the fact that, in the above example, the amounts are delivered (almost) immediately, we add the t (= current time) subscript: you pay an amount HC_t in home currency and you receive an amount FC_t of foreign currency.

By analogy to our exchange-of-cheques idea for a spot deal, then, we can picture a forward contract as an exchange of two promissory notes, with face values HC_T and FC_T, respectively:

Example 4.5

Suppose you sell forward USD 100,000 at EUR/USD 0.75 for December 31. (Note that the quote defines the euro as the HC.) Then

- you commit to deliver USD 100,000, which is similar to signing a promissory note (PN) with face value FC_T=USD 100,000 on Dec 31, and handing it over to the bank;
- the bank promises to pay you EUR 75,000, which is similar to giving you a signed PN with face value HC_T=EUR 75,000 for that date.

Intimately linked to the exchange markets are the money markets for the home and foreign country, that is, the markets for short-term deposits and loans. A home-currency deposit of GBP 1m "spot"[3] for one year at 4 percent means that you pay an amount of GBP 1m to the bank now, and the bank pays you an amount GBP 1.04m at time T. This is similar to handing over the spot money amount of HC_t=1m in return for a PN with face value HC_T=1.04m. Likewise, if you borrow GBP 10m at 6 percent over one year, this is tantamount to you receiving a cheque with face value HC_t=GBP 10m in return for a promissory note with face value HC_T=GBP 10.6m.

Graphical Representation of Chains of Transactions: an Example

For the remainder of this section, we take the Chilean Peso (CLP) as our home currency and the Norvegian Crown (NOK) as the foreign one. Suppose the spot rate is S_t=CLP/NOK 100, the four-year forward rate $F_{t,T}$=CLP/NOK 110, the CLP four-year risk-free rate of return is $r_{t,T}$=21 percent effective, and the NOK one equals $r^*_{t,T}$=10 percent. Very often we will discuss sequences of deals, or combinations of deals. Consider, for example, an Chilean investor who has CLP 100,000 to invest. He goes for a NOK deposit "swapped into CLP", that is, a NOK deposit combined with a spot purchase and a forward sale. Let us see what the final outcome is:

Example 4.6

The investor converts his CLP 100,000 into an amount NOK_t, deposits these for four years, and sells forward the proceeds NOK_T in order to obtain a risk-free amount of CLP four years from now. The outcome is computed as follows:

1. Buy spot NOK: the input given to the bank is CLP 100,000, so the output of the spot deal, received from the bank, is $100,000 \times 1/100 = 1,000$ Crown.
2. Invest these NOK at 10 percent: the input into the money market operation is NOK_t = 1,000, so after four years you will receive from the bank an output equal to $1,000 \times 1.10 = 1,100$ Crowns.

[3] A spot deposit or loan starts the second working day. For one-day deposits, one can also define the starting date as today ("overnight"), or tomorrow ("tomorrow/next"), but this must be made explicit, then. In all our examples, the deals are spot—the default option in real life, too.

3. This future NOK outcome is already being sold forward at t; that is, right now you immediately cover or hedge the NOK deposit in the forward market so as to make its time-T value risk free rather than contingent on the time-T spot rate. The input for this transaction is NOK_T = 1,100, and the output in CLP at time T will be $1,100 \times 110 = 121,000$.

There is nothing difficult about this, except perhaps that by the time you finish reading Step 3 you've already half forgotten the previous steps. We need a way to make clear at one glance what this deal is about, how it relates to other deals and what the alternatives are. One step in the right direction is to adopt a notation like

$$HC_t = 100,000 \overbrace{\rightarrow}^{\substack{\text{buy spot:} \\ \times 1/100}} FC_t = 1,000 \overbrace{\rightarrow}^{\substack{\text{deposit:} \\ \times 1.10}} FC_T = 1,100 \overbrace{\rightarrow}^{\substack{\text{sell frwd:} \\ \times 110}} HC_T = 121,000$$

So the arrows show how you go from a spot CLP position into a spot NOK one (the spot deal), and so on. We can further improve upon this by arranging the amounts in a diagram, where each kind of position has a fixed location. There are four kinds of money in play: foreign and domestic, each coming in a day-t and a day-T version. Let's show these on a diagram, with HC on the left and FC on the right, and with time t on top and time T below. Figure 4.1 shows the result for the above example.

We can now generalize. Suppose the spot rate is still CLP/NOK 100, the four-year forward rate CLP/NOK 110, the CLP risk-free is 21 percent effective, and the NOK one 10 percent. The diagram in Figure 4.2 summarises all transactions open to the treasurer. It is to be read as follows:

Figure 4.1: **Spot/Forward/Money Market Diagram: Example 4.6**

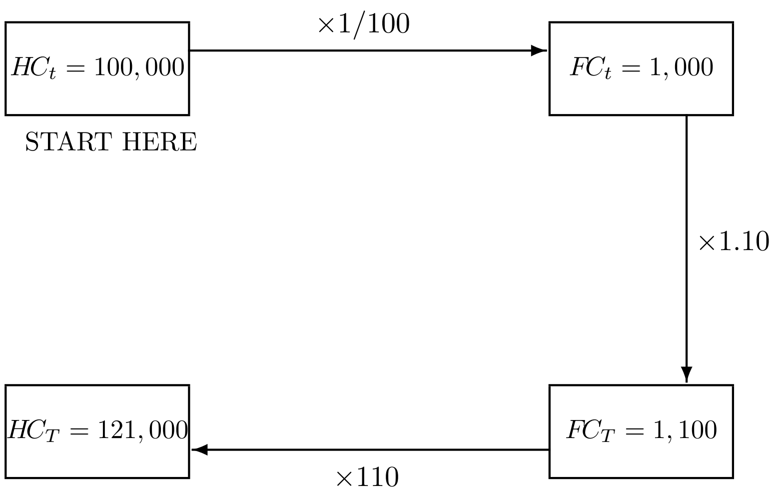

Figure 4.2: **Spot/Forward/Money Market Diagram: the general picture**

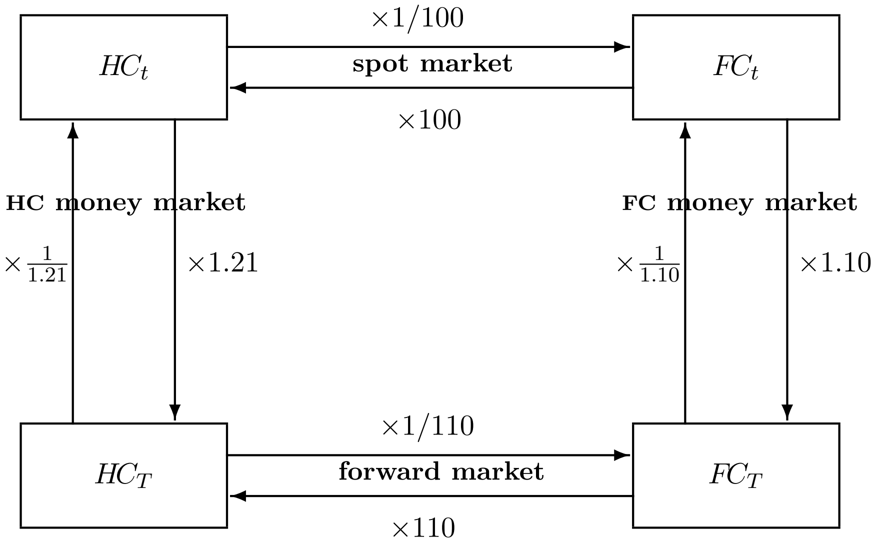

- Any t-subscripted symbol HC_t (FC_t) refers to an *amount* of spot money; and any T-subscripted symbol HC_T (FC_T) refers to a T-dated known amount of money, e.g. promised under a PN, A/P, A/R or deposit, or forward contract.

- Any possible transaction (spot or forward sale or purchase; home or foreign money-market deal) is shown as an arrow. A transaction is characterized by two numbers: (a) your position before the transaction, an input amount you surrender to the bank, and (b) your position after the transaction, the output amount you receive from the bank. The arrow starts from the (a) part and ends in the (b) part. For example,

  - a move $HC_t \rightarrow FC_t$ refers to *buying* FC—spot (see "t")
  - a move $FC_T \rightarrow HC_T$ refers to *selling* FC—forward (see "T")
  - a move $HC_t \rightarrow HC_T$ refers to *investing* or *lending* HC
  - a move $FC_T \rightarrow FC_t$ refers to *borrowing* against a FC income—*e.g.* discounting a FC PN.

- Next to each arrow we write the factor by which its "input" amount has to be multiplied to compute the "output" amount. Again: "input" is what you give to the bank (at either t or T), "output" is what you receive from it.

The General Spot/Forward/Money Market Diagram

To use the diagram, first identify the starting position. This is where you have money right now—like FC_T (: a customer will pay you FC in future, or a deposit will expire). Then determine the desired end point, like HC_T (: you want future HC instead; that is, you want to eliminate the exchange risk). Third, determine by which route you want to go from START to END. Lastly, follow the chosen route, sequentially multiplying the starting amount by all the numbers you see along the path.

Example 4.7

In Example 4.6, the path is $HC_t \to FC_t \to FC_T \to HC_T$, and the end outcome, starting from $HC_t = 100,000$ is immediately computed as

$$HC_T = 100,000 \times \frac{1}{100} \times 1.10 \times 110 = 121,000. \tag{4.5}$$

The alert reader will already have noted that this is a synthetic HC deposit, constructed out of a FC deposit and a swap, and that (here) it has exactly the same return as the direct solution. Indeed, the alternative route, $HC_t \to HC_T$, yields $100,000 \times 1.21 = 121,000$. (In *im*perfect markets this equivalence of both paths will no longer be generally true, as we shall see in the next chapter.)

Example 4.8

Suppose that a customer of yours will pay NOK 6.5m at time T, four years from now, but you need cash Pesos to pay your suppliers and workers. You decide to sell forward, and take out a CLP loan with a time-T value that, including interest, exactly matches the proceeds of the forward sale. How much can you borrow on the basis of this invoice without taking any exchange risk?

The path chosen is $FC_T(=\text{NOK}6,500,000) \to HC_T \to HC_t$, and it yields

$$6.5m \times 110 \times \frac{1}{1.21} = \text{CLP}\,590,909,090.91. \tag{4.6}$$

The clever reader will again eagerly point out that there is an alternative: borrow NOK against the future inflow (that is, borrow such that the loan *cum* interest is serviced by the NOK inflow), and convert the proceeds of the loan into CLP. Again, in our assumedly perfect market, the outcome is identical: $6.5m/1.10 \times 100 =$ CLP $590,909,090.91$. Thus, the diagram allows us to quickly understand the purpose, and see the outcome of, a sequence of transactions. It also shows there are always two routes that lead from a given starting point to a given end point—a useful insight for shopping-around purposes. The advantage of using the diagram will be even more marked when we add bid-ask spreads in all markets (next chapter) or when we study forward forwards or forward rate agreements and their relationship to forward contracts (Appendix 4.7),[4] or when we explain forward forward swaps (Chapter 5).

4.3 The Law of One Price and Covered Interest Parity

The sequences of transactions that can be undertaken in the exchange and money markets, as summarized in Figure 4.2, can be classified into two types.

1. You could do a sequence of transactions that forms a roundtrip. In terms of Figure 4.2, a roundtrip means that you start in a particular box, and then make four transactions that bring you back to the starting point. For example, you may consider the sequence $HC_T \to HC_t \to FC_t \to FC_T \to HC_T$. In terms of the underlying transactions, this means that you borrow CLP, convert the proceeds of the CLP loan into NOK, and invest these NOK; the proceeds of the investment are then immediately sold forward, back into CLP. The question that interests you is whether the CLP proceeds of the forward sale are more than enough to pay off the original CLP loan. If so, you have identified a way to make a sure profit without using any of your own capital. Thus, the idea behind a round-trip transaction is *arbitrage*, as defined in Chapter 3.
2. Alternatively, you could consider a sequence of transactions where you end up in a box that is not the same as the box from which you start. The two examples 4.7 and 4.8 describe such non-round-trip sequences. Trips like that have an economic rationale. In the first example, for instance, the investor wants to invest CLP, and the question here is whether the swapped NOK investment ($CLP_t \to NOK_t \to NOK_T \to CLP_T$) yields more than a direct CLP investment ($CLP_t \to CLP_T$). Using the terminology of Chapter 3, this would be an example of *shopping around* for the best alternative.

In what follows, we want to establish the following two key results:

1. To rule out arbitrage in perfect markets, the following equality must hold:

$$F_{t,T} = S_t \, \frac{1 + r_{t,T}}{1 + r^*_{t,T}}. \tag{4.7}$$

   [In *im*perfect markets, this sharp equality will be watered down to a zone of admissible values, but the zone is quite narrow.]

[4] Forward Forwards and forward rate agreements (FRAs) are contracts that fix the interest rate for a deposit or loan that will be made (say) six months from now, for (say) three months. This can be viewed as a six-month forward deal on a (then) three-month interest rate. See the Appendix on forward interest rates.

Figure 4.3: **Spot/Forward/Money Market Diagram: Arbitrage Computations**

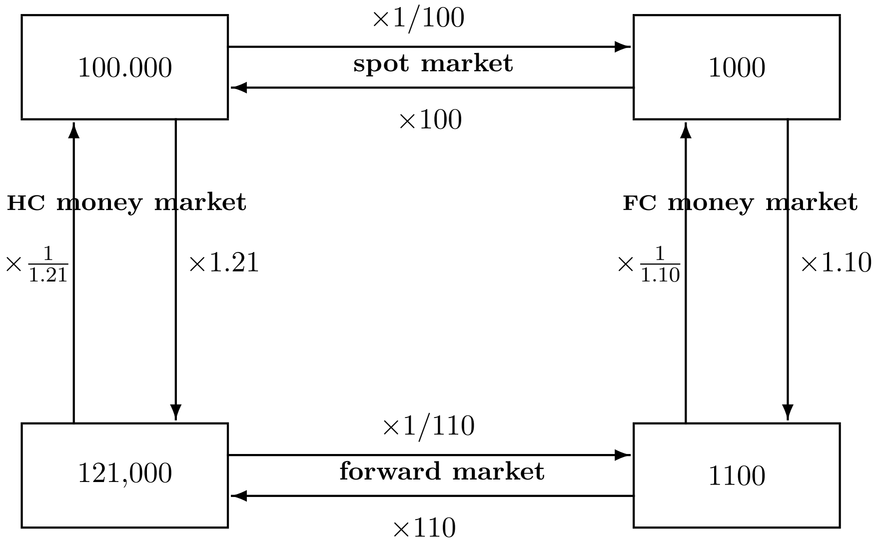

2. If equation (4.7) holds, shopping-around computations are a waste of time since the two routes that lead from a given initial position A to a desired end position B produce exactly the same. Stated positively, shopping around can (and will) be useful only because of imperfections.

4.3.1 Arbitrage and Covered Interest Parity

In this section, we use an arbitrage argument to verify equation (4.7), a relationship called the *Covered Interest Parity* (CIP) Theorem. The theorem is evidently satisfied in our example:

$$110 = F_{t,T} = S_t \frac{1 + r_{t,T}}{1 + r^*_{t,T}} = 100 \frac{1.21}{1.10} = 110. \tag{4.8}$$

Arbitrage, we know, means full-circle roundtrips through the diagram. There are two ways to go around the entire diagram: clockwise, and counterclockwise. Follow the trips on Figure 4.3, where the symbols for amounts have been replaced by the specific numbers used in the numerical examples. We should not make any profit if the rate is 110, and we should make free money as soon as the rate does deviate.

Clockwise roundtrip The starting point of a roundtrip is evidently immaterial, but let's commence with a HC loan: this makes it eminently clear that no own capital is being used. Also the starting amount is immaterial, so let's pick an amount that produces conveniently round numbers all around: we write a PN with face value

$CLP_T = 121,000$. We discount this,[5] and convert the proceeds of this loan into Crowns, which are invested. At the same moment we already sell forward the future Crown balance. The final outcome is:

$$121,000 \times \frac{1}{1.21} \times \frac{1}{100} \times 1.10 \times 110 = 121,000. \tag{4.9}$$

So we break exactly even: the forward sale nets us exactly what we need to pay back the loan.

DoItYourself problem 4.1

Show, similarly, that also the counterclockwise roundtrip exactly breaks even. For your convenience, start by writing a PN with face value $NOK_T = 1100$. What is the path? What is the outcome?

What if $F_{t,T}$ = too low, say 109? If one price is too low relative to another price (or set of other prices), we can make money by *buying* at this too-low rate. The trip where we buy forward is the counterclockwise one. We start as before, except for the new price in the last step:

$$1100 \times \frac{1}{1.10} \times 100 \times 1.21 \times \frac{1}{109} = 1110.09 > 1100. \tag{4.10}$$

So the forward purchase nets us 1110.09 Pesos, 10.09 more than the 1100 we need to pay back loan.

DoItYourself problem 4.2 :

What if $F_{t,T}$ is too high, say 111? Indicate the path and calculate the arbitrage profit.

DoItYourself problem 4.3

To generalize these numerical results, we now start with PN's with face value 1, and replace all rates by their symbols. One no-arb condition is that the proceeds of the clockwise trip should not exceed the starting amount, unity. Explain how this leads to the following expression:

$$\frac{1}{1 + r_{t,T}} \times \frac{1}{S_t} \times (1 + r^*_{t,T}) \times F_{t,T} \leq 1. \tag{4.11}$$

[5] Discounting a PN or a T-bill or a trade bill not only means computing its PV; it often means borrowing against the claim. In practice, under such a loan the borrower would typically also cede the claim to the financier, as security. This lowers the lender's risk and makes the loan cheaper.

This produces an inequality constraint, $F_{t,T} \leq S_t \frac{1+r_{t,T}}{1+r^*_{t,T}}$. Write the no-arbitrage-profit condition for the counterclockwise trip and express it as another inequality constraint. Lastly, derive CIP.

4.3.2 Shopping Around (The Pointlessness of —)

The diagram in Figure 4.2 also tells us that any non-round-trip sequence of transactions can be routed two ways. For instance, you can go directly from CLP_t to CLP_T, or you can go via NOK_t and NOK_T. In two earlier examples, 4.7 and 4.8, we already illustrated our claim that, in perfect markets where CIP holds, both ways to implement a trip produce exactly the same outcome. It is simple to show that this holds for all of the ten other possible trips one could think of, in this diagram; but it would also be so tedious that we leave this as an exercise to any non-believer in the audience. It would also be a bit pointless, because in reality shopping around does matter. As we show in the next chapter, the route you choose for your trip may matter because of imperfections like bid-ask spreads, taxes (if asymmetric), information costs (if leading to inconsistent risk spreads asked by home and foreign banks), and legal subtleties associated with swaps.

4.3.3 Unfrequently Asked Questions on CIP

Before we move on to new challenges like the market value of a forward contract and the relation of the forward rate with expected future spot rates, a few crucial comments are in order. We first talk about causality, then about why pros always quote the swap rate rather than the outright, and lastly about taxes.

Covered Interest Parity and Causality

As we have seen, in perfect markets the forward rate is linked to the spot rate by pure arbitrage. Such an arbitrage argument, however, does not imply any causality. CIP is merely an application of the Law of One Price, and the statement that two perfect substitutes should have the same price does not tell us where that "one price" comes from. Stated differently, showing $F_{t,T}$ as the left-hand-side variable (as we did in Equation [4.7]), does not imply that the forward rate is a "dependent" variable, determined by the spot rate and the two interest rates. Rather, what Covered Interest Parity says is that the four variables (the spot rate, the forward rate, and the two interest rates) are determined jointly, and that the equilibrium outcome should satisfy Equation [4.7]. The fact that the spot market represents less than 50 percent of the total turnover likewise suggests that the forward market is not just an appendage to the spot market. Thus, it is impossible to say, either in theory or in practice, which is the tail and which is the dog, here.

Figure 4.4: **Appending Underlying Stories to the Variables in CIP**

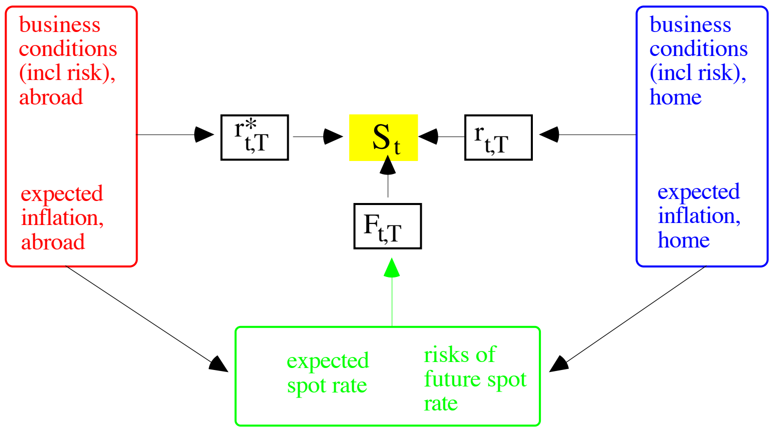

Although CIP itself does not say which term causes which, many economists and practitioners do have theories about one or more terms that appear in the Covered Interest Parity Theorem. One such theory is the Fisher equation, which says that interest rates reflect expected inflation and the real return that investors require. Another theory suggests that the forward rate reflects the market's expectation about the (unknown) future spot rate, $\tilde{S}_T$.[6] We shall argue in Section 4.4 that the latter theory is true in a risk-adjusted sense. In short, while there is no causality in CIP itself, one can append stories and theories to items in the formula. Then CIP becomes an ingredient in a richer economic model with causality relations galore—but S, F, r and r^* would all be endogenous, determined by outside forces and circumstances. Figure 4.4 outlines a plausible causal story of how interest rates and the forward rate are set and, together, imply the spot rate.

CIP and the Swap Rate

When the forward rate exceeds the spot rate, the foreign currency is said to be *at a premium*. Otherwise, the currency is *at a discount* ($F_{t,T} < S_t$), or *at par* ($F_{t,T} = S_t$). In this text, we often use the word premium irrespective of its sign; that is, we treat the discount as a negative premium. From [4.7], the sign of the premium uniquely

[6]We use a tilde ($\tilde{\cdot}$) above a symbol to indicate that the variable is random or uncertain.

depends on the sign of $r_{t,T} - r^*_{t,T}$:

$$
\begin{aligned}
[\text{swap rate}]_{t,T} &\stackrel{\text{def}}{=} F_{t,T} - S_t, \\
&= S_t \left[\frac{1 + r_{t,T}}{1 + r^*_{t,T}} - 1 \right], \\
&= S_t \left[\frac{1 + r_{t,T}}{1 + r^*_{t,T}} - \frac{1 + r^*_{t,T}}{1 + r^*_{t,T}} \right], \\
&= S_t \left[\frac{r_{t,T} - r^*_{t,T}}{1 + r^*_{t,T}} \right]; \qquad (4.12) \\
\Rightarrow \frac{\partial \cdot}{\partial S_t} &= \left[\frac{r_{t,T} - r^*_{t,T}}{1 + r^*_{t,T}} \right] \approx r_{t,T} - r^*_{t,T}. \qquad (4.13)
\end{aligned}
$$

Thus, a higher domestic return means that the forward rate is at a premium, and vice versa. To a close approximation (with low foreign interest rates and/or short maturities), the percentage swap rate even simply is the effective return differential.

To easily remember this, think of the following. If there would be a pronounced premium, we would tend to believe that this signals an expected appreciation for the foreign currency.[7] That is, the foreign currency is “strong”. But strong currencies are also associated with low interest rates: it’s the weak moneys that have to offer high rates to shore up their current value. In short, a positive forward premium goes together with a low interest rate because both are traditionally associated with a strong currency.

A second corollary from the CIP theorem is that, whenever the spot rate changes, all forward rates must change in lockstep. In old, pre-computer days, this meant quite a burden to traders/market makers, who would have to manually recompute all their forward quotes. Fortunately, traders soon noticed that the swap rate is relatively insensitive to changes in the spot rate. That is, when you quote a spot rate and a swap rate, then you make only a small error if you do not change the swap rate every time S changes.

Example 4.9

Let the *p.a.* simple interest rates be 4 and 3 percent (HC and FC, respectively). If S_t changes from 100 to 100.5—a huge change—the theoretical one-month forward changes too, and so does the swap rate, but the latter effect is minute:

[7] Empirically, the strength of a currency is predicted by the swap rate only in the case of pronounced premia. When interest are quite similar and expectations rather diffuse, as is typically the case among OECD mainstream countries, the effects risk premia and transaction costs appear to swamp any expectation effect. See Chapter 10.

	spot		forward	swap rate
level_0	100.0	$100.0 \times \frac{1.003333}{1.002500}$	$= 100.0831$	0.0831
level_1	100.5	$100.5 \times \frac{1.003333}{1.002500}$	$= 100.5835$	0.0835
change	0.5		0.5004	0.0004

The rule of thumb of not updating the swap rate all the time used to work reasonably well because, in olden days, interest rates were low[8] and rather similar across currencies (the gold standard, remember?), and maturities short. This makes the fraction on the right hand side of [4.12] a very small number. In addition, interest rates used to vary far less often than spot exchange rates. Nowadays, of course, computers make it very easy to adjust all rates simultaneously without creating arbitrage opportunities, so we no longer need the trick with the swap rates. But while the motivation for using swap rates is gone, the habit has stuck.

DoItYourself problem 4.4

Use the numbers of Example 4.9 to numerically evaluate the partial derivative in Equation [4.13],

$$\frac{\partial(F_{t,T} - S_t)}{\partial S_t} = \left[\frac{r_{t,T} - r^*_{t,T}}{1 + r^*_{t,T}}\right] \approx r_{t,T} - r^*_{t,T} \qquad .$$

Check whether this is a small number, when interest rates are low (and rather similar across currencies) and maturities short. (If so, it means that the swap rate hardly changes when the spot rate moves.) Also check that the analytical result matches the calculations in the Example.

We now bring up an issue we have been utterly silent about thus far: taxes.

CIP: Capital Gains *v* Interest Income, and Taxes

When comparing the direct and synthetic HC deposits, in Example 4.7, we ignored taxes. This, we now show, is fine as long as the tax law does not discriminate between interest income and capital gains.

[8] During the Napoleonic Wars, for instance, the UK issued perpetual (!) debt (the consolidated war debt, or *consol*) with an interest rate of 3.25 percent. Toward the end of the nineteenth century, Belgium issued perpetual debt with a 2.75 percent coupon (to pay off a Dutch toll on ships plying for Antwerp). Rates crept up in the inflationary 70s to, in some countries, 20 percent short-term or 15 percent long-term around 1982. They then fell slowly to quite low levels, as a result of falling inflation, lower government deficits and, in the first years of the 21st century, high uncertainty and a recession—the "flight for safety" effect.

Table 4.2: **HC and swapped FC investments with nondiscriminatory taxes**

	Invest CLP 100		Invest NOK 1 and hedge	
initial investment		100	$1 \times 100 =$	100
final value	$100 \times 1.21 =$	121	$[1 \times 1.10] \times 110 =$	121
income		21		21
interest		21	$[1 \times 0.10] \times 110 =$	11
capgain		0	$110 - 100 =$	10
taxable		21		21
tax (33.33 %)		7		7
after-tax income		14		14

The first point you should be aware of is that, by going for a swapped FC deposit instead of a HC one, the total return is in principle unaffected but the relative weight of the interest and capital-gain components is changed. Consider our Chilean investor who compares an investment in NOK to one in CLP. Given the spot rate of 100, we consider investments of 100 CLP or 1 NOK. In Table 4.2 you see that the CLP investment yields interest income only, while the NOK deposit earns interest (10 pence, exchanged at the forward rate 110) and a capital gain (you buy the principal at 100, and sell later at 110). But in both cases, total income is 21. (This, indeed, is the origin of the name CIP: the return, covered, is the same.)

DoItYourself problem 4.5

Verify that the expression below follows almost immediately from CIP, Equation [4.7]:

$$F_{t,T}\, r^*_{t,T} + (F_{t,T} - S_t) = S_t\, r_{t,T}. \tag{4.14}$$

Then trace each symbol in the formula to the numbers we used in the numerical example. Identify the interest on the Peso and Crown deposits, and the capital gain or loss.

So we know that total pre-tax income is the same in both cases. If all income is equally taxable, the tax is the same too, and so must be the aftertax income. It also follows that if, because of *e.g.* spreads, there is a small advantage to, say, the Peso investment, then taxes will reduce the gain but not eliminate it. That is, if Pesos would yield more before taxes, then they would also yield more after taxes.

In most countries, corporate taxes are neutral between interest income and capital gains, especially short-term capital gains. But there are exceptions. The UK used to treat capital gains on FC loans differently from capital losses and interest received. Under personal taxation, taxation of capital gains is far from universal, and/or long-term capital gains often receive beneficial treatment. In cases like this, the ranking of outcomes on the basis of after-tax returns could be very different from the ranking on pre-tax outcomes. Beware!

4.4 The Market Value of an Outstanding Forward Contract

In this and the next section, we discuss the market value of a forward contract at its inception, during its life, and at expiration. As is the case for any asset or portfolio, the market value of a forward contract is the price at which it can be bought or sold in a normally-functioning market. The focus, in this section, is on the value of a forward contract that was written in the past but that has not yet matured. For instance, one year ago (at time t_0), we may have bought a five-year forward contract for NOK at $F_{t_0,T}$ = CLP/NOK 115. This means that we now have an outstanding four-year contract, initiated at the rate of CLP/NOK 115. This outstanding contract differs from a newly signed four-year forward purchase because the latter would have been initiated at the now-prevailing four-year forward rate, CLP/NOK 110. The question then is, how should we value the outstanding forward contract?

This value may be relevant for a number of reasons. At the theoretical level, the market value of a forward contract comes in quite handy in the theory of options, as we shall see later on. In day-to-day business, the value of an outstanding contract can be relevant in, for example, the following circumstances:

- If we want to negotiate early settlement of the contract, for instance to stop losses on a speculative position, or because the underlying position that was being hedged has disappeared.
- If there is default and the injured party wants to file a claim.
- If a firm wishes to "mark to market" the book value of its foreign-exchange positions in its financial reports.

4.4.1 A general formula

Let us agree that, unless otherwise specified, "a contract" refers to a forward purchase of one unit of foreign currency. (This is the standard convention in futures markets.) Today, at time t, we are considering a contract that was signed in the past, at time t_0, for delivery of one unit of foreign currency to you at T, against payment of the initially agreed-upon forward rate, $F_{t_0,T}$. Recall the convention that we have adopted for indicating time: the current date is always denoted by t, the initiation date by t_0, the future (maturity) date by T, and we have, of course, $t_0 \leq t \leq T$.

The way to value an outstanding contract is to interpret it as a simple portfolio that contains a FC-denominated PN with face value 1 as an asset, and a HC-denominated PN with face value $F_{t_0,T}$ as a liability. Valuing a HC PN is easy: just discount the face value at the risk-free rate. For the FC PN, we first compute its PV in FC (by discounting at r^*), and then translate this FC value into HC via the spot price:

Example 4.10

Consider a contract that has 4 years to go, signed in the past at a historic forward price of 115. What is the market value if $S_t = 100$, $r_{t,T} = 21\%$, $r^*_{t,T} = 10\%$?

- The asset leg is like holding a PN of FC 1, now worth PV* = 1/1.10 = 0.90909 NOK and, therefore, 0.9090909 × 100 = 90.909 CLP.
- The liability leg is like having written a PN of HC 115, now worth CLP 115/1.21 = 95.041.
- The net value now is, therefore, CLP 90.909 – 95.041 = –4.132

The generalisation is as follows:

$$\begin{bmatrix}\text{Market value of}\\ \text{forward purchase}\\ \text{at } F_{t_0,T}\end{bmatrix} = \underbrace{\overbrace{\frac{1}{1+r^*_{t,T}}}^{\text{PV* of asset, FC 1}} \times S_t}_{\text{translated value of FC asset}} - \underbrace{\frac{F_{t_0,T}}{1+r_{t,T}}}_{\text{PV of HC liability}} . \tag{4.15}$$

There is a slightly different version that is occasionally more useful: the value is the discounted difference between the current and the historic forward rates. To find this version, multiply and divide the first term on the right of [4.15] by $(1 + r_{t,T})$, and use CIP:

$$\begin{aligned}\begin{bmatrix}\text{Market value of}\\ \text{forward purchase}\\ \text{at } F_{t_0,T}\end{bmatrix} &= \frac{1}{1+r_{t,T}} \underbrace{\frac{1+r_{t,T}}{1+r^*_{t,T}} S_t}_{=F_{t,T},\ \text{(CIP)}} - \frac{F_{t_0,T}}{1+r_{t,T}} \\ &= \frac{F_{t,T} - F_{t_0,T}}{1+r_{t,T}} \end{aligned} \tag{4.16}$$

Example 4.11

Go back to Example 4.10. Knowing that the current forward rate is 110, we immediately find a value of (110 – 115)/1.21 = –4.132 CLP for a contract with historic rate 115.

One way to interpret this variant is to note that, relative to a new contract, we're overpaying by CLP 5: last year we committed to paying 115, while we would have gotten away with 110 if we had signed right now. This "loss", however, is dated 4 years from now, so its PV is discounted at the risk-free rate.

The sceptical reader may object that this "loss" is very fleeting: its value changes every second; how comes, then, that we can discount at the risk-free rate? One answer is that the value changes continuously because interest rates and (especially) the spot rate are in constant motion, But that does not invalidate the claim that we can always value each PN using the risk-free rates and the spot exchange rate prevailing at that moment. Relatedly, the future loss relative to market conditions at t can effectively be locked in at no cost, by selling forward for the same date:

Example 4.12

Consider a contract that has four years to go, signed in the past at a historic forward price of 115, for speculative purposes. Right now you see there is a loss, and you want to close out to avoid any further red ink. One way is to sell forward HC 1 at the current forward rate, 110. On the common expiry date of old and new contract we then just net the loss of 115–110:

	HC flows at T	FC flows at T
old contract: buy at $F_{t_0,T}$=115	–115	1
new contract: sell at $F_{t_0,T}$=110	110	–1
net flow	–5	0

But because this loss is realized within four years only, its PV is found by discounting. Discounting can be at the risk-free rate since, as we see, the locked-in loss is risk free.

We can now use the result in Equation [4.15] to determine the value of a forward contract in two special cases: at its inception and at maturity.

4.4.2 Corollary 1: The Value of a Forward Contract at Expiration

At its expiration time, the market value of a purchase contract equals the difference between the spot rate that prevails at time T—the value of what you get—and the forward rate $F_{t_0,T}$ that you agreed to pay:

$$\begin{bmatrix}\text{Expiration value of}\\ \text{a forward contract}\\ \text{with rate } F_{t_0,T}\end{bmatrix} = S_T - F_{t_0,T}. \tag{4.17}$$

Equation [4.17] can be derived formally from Equation [4.15], using the fact that the effective return on a deposit or loan with zero time to maturity is zero (that is, $r_{T,T} = 0 = r^*_{T,T}$. The result in [4.17] is quite obvious, as the following example shows:

Example 4.13

- You bought forward, at time t_0, one NOK at CLP/NOK 115. At expiry, T, the NOK spot rate turns out to be CLP/NOK 123, so you pay 115 for something you can immediately re-sell at 123. The net value is, therefore, 123-115=8.

- Idem, except that S_T turns out to be CLP/NOK 110. You have to pay 115 for something worth only 110. The net value is, therefore, 110-115=–5: you would be willing to *pay* 5 to get out of this contract.

The value of a unit forward sale contract is, of course, just the negative of the value of the forward purchase: forward deals are zero-sum games. The seller wins if the spot value turns out to be below the contracted forward price, and loses if the spot value turns out to be above. Figure 4.5 pictures the formulas, with smileys and frownies indicating the positive and negative parts.

Equation [4.17] can be used to formally show how hedging works. Suppose that you have to pay one unit of foreign currency at some future time T. The foreign currency debt is risky because the cash flow at time T, in home currency, will be equal to minus the future spot rate—and, at time t, this future spot rate is uncertain, a characteristic we stress by adding a tilde ($\sim$) over the variable. By adding a forward purchase, the combined cash flow becomes risk free, as the bit of arcane math shows, below:

$$\begin{array}{rr} \text{Cash flow from amortizing the debt at expiration:} & -\tilde{S}_T \\ \text{Value of the forward purchase at expiration:} & \underline{\tilde{S}_T - F_{t_0,T}} \\ \text{Combined cash flow:} & -F_{t_0,T}. \end{array} \tag{4.18}$$

Putting this into words, we say that hedging the foreign-currency debt with a forward purchase transforms the risky debt into a risk-free debt, with a known outflow

Figure 4.5: **The Value of a Forward Purchase or Sales Contract at Expiry**

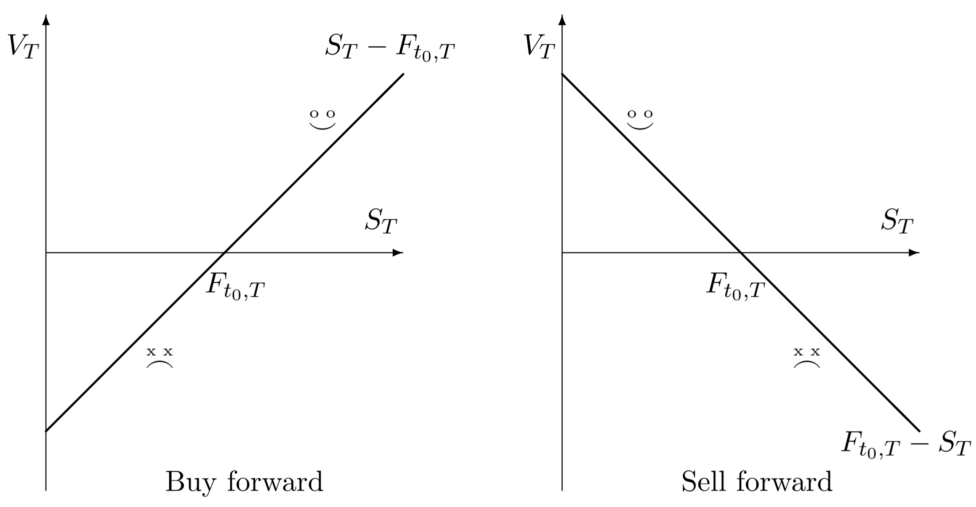

$-F_{t_0,T}$. We shall use this result repeatedly in Chapter 5 (on uses of forward contracts), in Chapter 9 when we discuss option pricing, and in Chapters 13 where we analyse exposure and risk management.

Make sure you realize that the hedged liability may make you worse off, *ex post*, than the unhedged one. Buying at a pre-set rate $F_{t,T}$ gives that great warm feeling, *ex post*, if the spot rate S_T turns out to be quite high; but it hurts if the spot rate turns out to be quite cheap. The same conclusion was already implicit in [4.17]: the value of the contract at expiry can be either sign. This rises the question whether hedging is really so good as it is sometimes cracked up to be. We return to the economics of hedging in Chapter 12.

4.4.3 Corollary 2: The Value of a Forward Contract at Inception

The value at expiry, above, probably was so obvious that it is, in a way, just a means of proving that the general valuation formula [4.15] makes sense. The same holds for the next special case: the value at inception, *i.e.* the time the contract is initiated or signed. At inception, the market value *must* be zero. We know this because (a) when we sign a forward contract, we have to pay nothing; and (b) hard-nosed bankers would never give away a positive-value contract for free, nor accept a negative-value contract at a zero price. To show the (initial) zero-value property formally, we use the general value formula [4.16] and consider the special case where $t_0 = t$, implying that $F_{t_0,T} = F_{t,T}$ (That is, the contract we are valuing is new.) Obviously,

$$\begin{bmatrix}\text{Initial value of a}\\ \text{forward contract}\\ \text{with rate } F_{t,T}\end{bmatrix} = \frac{F_{t,T} - F_{t,T}}{1 + r_{t,T}} = 0. \tag{4.19}$$

The value of a forward contract is zero at the moment it is signed because the contract can be replicated at zero cost. Notably, if a bank tried to charge you money for a contract at the equilibrium (Covered Interest Parity) forward rate, you would refuse, and create a synthetic forward contract through the spot and money markets:

Example 4.14

Let $S_t = 100$, $r^*_{t,T} = 0.10$, $r_{t,T} = 0.21$, $F_{t,T} = 110$; but a bank wants to charge you a commission of 3 for a forward purchase. You would shrug dismissively and immediately construct a synthetic forward contract at 110 at a zero cost:

- write a PN *ad* HC 110, discount it;
- convert the proceeds, 110/1.21 = 90.909090, into FC: you get FC 0.90909.
- invest at 10 percent, to get HC 1 at T.

Thus, you can replicate a forward purchase contract under which your payment at T amounts to 110, just like in the genuine, direct forward contract, but it does not cost you anything now.

4.4.4 Corollary 3: The Forward Rate and the Risk-Adjusted Expected Future Spot Rate

The zero-value property of forward contracts discussed above has another, and quite fundamental, interpretation. Suppose that the CLP/NOK four-year forward rate equals 110, implying that you can exchange one future NOK for 110 future CLP and *vice versa* without any up-front cash flow. This must mean that the market perceives these amounts as being *equivalent* (that is, having the same value). If this were not so, there would have been an up-front compensation to make up for the difference in value.

Since any forward contract has a zero value, the present values of CLP 1 four years and NOK 110 four years must be equal anywhere; that is, the equivalence of these amounts holds for any investor or hedger anywhere. However, the equivalence property takes on a special meaning if we pick the CLP (which is the currency in which our forward rate is expressed) as the home currency: in that particular numéraire, the CLP amount is risk-free, or certain. In terms of CLP, we can write the equal-value property as:

$$\mathrm{PV}_t(\tilde{S}_T) = \mathrm{PV}_t(F_{t,T}), \tag{4.20}$$

where $\mathrm{PV}_t(.)$ is the present-value operator. In a way, equation [4.20] is just the zero-value property: the present value of the uncertain future cash inflow $\tilde{S}_T$ generated by the contract cancels out against the PV of the known future outflow, $F_{t,T}$. We can lose or gain, but these prospects balance out in present-value terms, from our time-t viewpoint. But the related, second interpretation stems from the fact that in home currency, the forward price on the right-hand side of Equation [4.20] is a risk-free, known number whereas the future spot rate on the left is uncertain. That is, at time t an amount of $F_{t,T}$ Pesos payable at T is not just *equivalent* to one unit of foreign currency payable at T; this amount of future home currency is also a *certain*, risk-free amount. For this reason, we shall say that in home currency, the forward rate is the time-t *certainty equivalent* of the future spot rate, $\tilde{S}_T$.

Example 4.15

In our earlier CLP/NOK examples, the certainty equivalent of one Norvegian Crown four years out is CLP 110. You can offer the market a sure CLP 110 at T and get one Crown (with risky value $\tilde{S}_T$) in return; but equally well you can offer the market one Crown (with risky value $\tilde{S}_T$) and get a sure CLP 110 in return.

The notion of the certainty equivalent deserves some elaboration. Many introductory finance books discuss the concept of an investor's subjective certainty equivalent of a risky income. This is defined as the single known amount of income that is equally attractive as the entire risky distribution.

Example 4.16

Suppose that you are indifferent between, on the one hand, a lottery ticket that pays out with equal probabilities either USD 100m or nothing, and on the other hand, a sure USD 35m. Then your personal certainty equivalent of the risky lottery is USD 35m. You are indifferent between 35m for sure and the risky cash flow from the lottery.

Another way of saying this is that, when valuing the lottery ticket, you have marked down its expected value, USD 50m, by USD 15, because the lottery is risky. Thus, we can conclude that your personal certainty equivalent, USD 35m, is the expected value of the lottery ticket corrected for risk.[9]

In the example, the risk-adjustment is quite subjective. A *market* certainty equivalent, by analogy, is the single known amount that the market considers to be as valuable as the entire risky distribution. And market certainty equivalents are, of course, what matter if we want to price assets, or if we want to make managerial decisions that maximize the market value of the firm. We have just argued that the (CLP) market certainty equivalent of the future CLP/NOK spot rate must be the current CLP/NOK forward rate. Stated differently, the market's time-t expectation of the time-T CLP/NOK spot rate, corrected for risk, is revealed in the CLP/NOK forward rate, $F_{t,T}$. Let's express this formally as:

$$\mathrm{CEQ}_t(\tilde{S}_T) = F_{t,T}, \tag{4.21}$$

where $\mathrm{CEQ}_t(.)$ is called the *certainty equivalent* operator.

A certainty equivalent operator is similar to an ordinary expectations operator, $\mathrm{E}_t(.)$, except that it is a risk-adjusted expectation rather than an ordinary expected value. (There are good theories as to how the risk-adjusted and the "physical" densities are related, but they are beyond the scope of this text.) Like $\mathrm{E}_t(.)$, $\mathrm{CEQ}_t(.)$ is also a conditional expectation, that is, the best possible forecast given the information available at time t. We use a t subscript to emphasize this link with the information available at time t.

To make the market's risk-adjustment a bit less abstract, assume the CAPM holds. Then we could work out the left-hand side of [4.20] in the standard way: the PV of a risky cashflow $\tilde{S}_T$ equals its expectation, discounted at the risk-adjusted rate. The risk-adjusted discount rate, in turn, consists of the risk-free rate plus a risk premium $RP_{t,T}(\beta_S)$ which depends on market circumstances and the risk of the asset to be priced, β_S. Working out the right-hand side of [4.20] is straightforward: the PV of a risk-free flow F is F discounted at the risk-free rate r. Thus, we can

[9] When we say that investors are risk-avert, we mean they do not like symmetric risk for their entire wealth. The amounts in the example are so huge that they would represent almost the entire wealth of most readers; so in that case, risk aversion guarantees that the risk-adjustment is downward. But for small investments with, for instance, lots of right skewness, one observes upward adjustments: real-world lottery players, for instance, are willing to pay more than the expected value because, when stakes are small, right-skewness can give quite a kick.

flesh out [4.20] into

$$\frac{\mathrm{E}_t(\tilde{S}_T)}{1+r_{t,T}+RP_{t,T}(\beta_S)} = \frac{F_{t,T}}{1+r_{t,T}}. \tag{4.22}$$

After a minor rearrangement (line 1, below)we can then use the notation CEQ as in [4.21], to conclude that

$$\begin{aligned} F_{t,T} &= \mathrm{E}_t(\tilde{S}_T)\frac{1+r_{t,T}}{1+r_{t,T}+RP_{t,T}} \\ \Rightarrow \mathrm{CEQ}_t(\tilde{S}_T) &= \mathrm{E}_t(\tilde{S}_T)\frac{1+r_{t,T}}{1+r_{t,T}+RP_{t,T}} \qquad (4.23) \\ &\approx \mathrm{E}_t(\tilde{S}_T)\,\frac{1}{1+RP_{t,T}}. \qquad (4.24) \end{aligned}$$

The last line is only an approximation of the true relation [4.23]. We merely add it to show why the fraction on the right-hand side of [4.23] is called the risk adjustment.

Example 4.17

Suppose your finance professor offers you a 1-percent share in the next-year royalties from his finance textbook, with an expected value, next year, of USD 3,450,000. Given the high risk (☺), the market would discount this at 10 percent—3 risk-free plus a 7 risk premium. The CEQ would be

$$\mathrm{CEQ} = 3,450,000\frac{1.03}{1.10} = 3,450,000 \times 0.936363636 = 3,090,000. \tag{4.25}$$

Thus, the market would be indifferent between this proposition and USD 3,090,000 for sure. You could unload either of these in the market at a common PV,

$$\frac{3,450,000}{1.10} = 3,000,000 = \frac{3,090,000}{1.03}. \tag{4.26}$$

The risk-adjusted expected value plays a crucial role in the theory of international finance. As we shall see in the remainder of this chapter and later in the book, the risk-adjusted expectation has many important implications for asset pricing as well as for corporate financial decisions.

4.4.5 Implications for Spot Values; the Role of Interest Rates

In principle, we can see the spot value as the expected future value of the investment—including interest earned—corrected for risk and then discounted at the appropriate risk-free rate. In this subsection we consider the role of interest rates and changes therein, hoping to clear up any confusion that might exist in your mind. Notably, we have noted that a forward discount, *i.e.* a relatively high foreign riskfree rate, signals a weak currency. Yet we see central banks increase interest rates when their currency is under pressure, and the result often is an appreciation of the spot value. How can increasing the interest rate, a sign of weakness, strengthen the currency?

The relation to watch is, familiarly,

$$S_t = \frac{\mathrm{CEQ}_t(\tilde{S}_T)(1 + r^*_{t,T})}{1 + r_{t,T}}. \tag{4.27}$$

We also need to be clear about what is changing, here, and what is held constant. Let's use an example to guide our thoughts.

Example 4.18

Assume that the CAD (home currency) and GBP risk-free interest rates, $r_{t,T}$ and $r^*_{t,T}$, are both equal to 5 percent *p.a.* Then, from Equation [4.27], initially no change in S is expected, after risk adjustments: the spot rate is set equal to the certainty equivalent future value. Now assume that bad news about the British (foreign) economy suddenly leads to a downward revision of the expected next-year spot rate from, say, CAD/GBP 2 to 1.9. From Equation [4.27], if interest rates remain unchanged, the current spot rate would immediately react by dropping from 2 to 1.9, too. Exchange rates, like any other financial price, anticipate the future.

Now if the Bank of England does not like this drop in the value of the GBP, it can prop up the current exchange rate by increasing the British interest rate. To do this, the UK interest rate will need to be increased from 5 percent to over 10.5 percent, so that S_t equals CAD/GBP 2 even though $\mathrm{CEQ}_t(\tilde{S}_T)$ equals 1.9:

$$\frac{1.9 \times 1.10526}{1.05} = 2. \tag{4.28}$$

Thus, the higher UK interest rate does strengthen the current GBP spot rate, all else being equal.

But this still means that the currency is weak, in the sense that the GBP is still expected to drop towards 1.9, after risk-adjustment, in the future. Actually, in this story the pound strengthens now so that it can become weak afterwards. So there is no contradiction, since "strengthening" has to do with the immediate spot rate (which perks up as soon as the UK interest rate is raised, holding constant the CEQ), while "weakening" refers to the expected movements in the future.

A second comment is that, in the example, the interest-rate hike merely postpones the fall of the pound to a risk-adjusted 1.9. In this respect, however, this partial analysis may be incomplete, because a change in interest rates may also affect expectations. For instance, if the market believes that an increase in the British interest rate also heralds a stricter monetary policy, this would increase the expected future spot rate, and reinforce the effect of the higher foreign interest rate. Thus, the BoE would get away with a lower rise in the UK interest rate than in the first version of the story.

Of course, if expectations change in the opposite direction, the current spot rate may decrease even when the foreign interest rate is increased. For example, if the foreign interest rate rises by a smaller amount than was expected by the market, this may then lead to a downward revision of the expected future exchange rate and, ultimately, a drop in the spot value.

Example 4.19

Suppose that the current interest rates are equal to 5 percent p.a. in both Canada (the home country) and the UK, and the current and expected exchange rate are CAD/GBP 2. The Bank of England now increases its interest rate to 5.025 percent p.a. in an attempt to stem further rises in UK inflation. It is quite possible that this increase in interest rates is interpreted by the market as a negative signal about the future state of the UK economy (the BoE wants to slow things down) or as insufficient to stop inflation. So the market may revise expectations about the CAD/GBP exchange rate from 2 to 1.95. Thus, the change in the interest rate is insufficient to match the drop in the expected exchange rate. Instead of appreciating, the current exchange rate drops to CAD/GBP $1.95 \times 1.0525/10,5 = 1.955$.

Note the difference between the two examples. In the first, there was a drop in expectations that was perfectly offset by the interest rate—for the time being, that is: the drop is just being postponed, by assumption. In the second example the interest rate change came first, and then led to a revision of expectations. So we need to be careful about expectations too when the role of interest rates is being discussed.

Let's return to more corporate-finance style issues:

4.4.6 Implications for the Valuation of Foreign-Currency Assets or Liabilities

The certainty-equivalent interpretation of the forward rate implies that, for the purpose of corporate decision making, one can use the forward rate to translate foreign-currency-denominated claims or liabilities into one's domestic currency without much ado. Indeed, identifying the true expectation and then correcting for risk would just be reinventing the wheel: the market has already done this for you, and has put the result upon the Reuters screen. This makes your life much more simple. Rather than having to tackle a valuation problem involving a risky cash flow—the left-hand side of Equation [4.22]—we can simply work with the right-hand side where the cash flow is risk free. With risk-free cash flows, it suffices to use the observable domestic risk-free rate for discounting purposes.

Example 4.20

If the domestic CLP risk-free return is 21 percent, effective for 4 years, and the 4-year forward rate is CLP/NOK 110, then the (risk-adjusted) economic value of a NOK 5,000 4-year zero-coupon bond can be found as

$$\frac{\text{NOK}\,5,000 \times \text{CLP/NOK}\,110}{1.21} = \text{CLP}\ 454,545.45, \tag{4.29}$$

without any fussing and worrying about expectations or risk premia.

As illustrated in the example, the expected spot rate is not needed in order to value this position, and discounting can be done at the risk-free rate of return. In contrast, if you had tried to value the position using the left-hand side of Equation [4.22], you would probably have had to discount the expected future spot rate at some risk-adjusted rate. Thus, the first problem would have been to estimate the expected future spot rate. Unlike the forward rate, this expectation is not provided in the newspaper or on the Reuters screens. Second, you would have had to use some asset-pricing theory like the international Capital Asset Pricing Model (CAPM) to calculate a risk-adjusted discount rate that we use on the left-hand side of Equation [4.22]. In this second step, you would run into problems of estimating the model parameters, not the mention the issue of whether the CAPM is an appropriate model. In short, the forward rate simplifies decision making considerably. We shall use this concept time and again throughout this text.

4.4.7 Implication for the Relevance of Hedging

In this mercifully short last section before the wrap-up, we briefly touch upon the implications of the zero initial value for the relevance of hedging, that is, using financial instruments to reduce or even entirely eliminate the impact of exchange rates on the cash flow. Forward contracts are a prime instrument for this purpose: if one contractually fixes the rates at which future exchanges will be made, then the future spot rate no longer affects your bank account—at least not for those transactions.

The zero-value property has been invoked by some (including me, when *very* young) as implying that such hedging does not add value, or more precisely that any value effects must stem from market imperfections. This is wrong, but it took me some time to figure out exactly what was wrong.

The argument views the firm as a bunch of cash-flow-generating activities, to which a hedge is added. The cash flow triggered by the hedge is some positive or negative multiple of $\tilde{S}_T - F_{t,T}$, and its PV is zero the moment the hedge contract is signed. True, it's value will become non-zero one instant later, but we have no clue whether this new value will be positive or negative; so our knowledge that the zero-value property is short-lived is of no use for hedging decisions. But does zero initial value mean that the hedge is (literally) worthless? There can be, and will be, a value effect if the firm's other cash flows are affected. For instance, the chances that adverse currency movements wipe out so much capital that R&D investments must be cut, or that banks increase their risk spreads on loans, or customers desert the company, or the best employees leave like rats from a sinking ship—the chances that all these bad things happen should be lower, after hedging. Perhaps the firm is so well off that the probability of painful bad luck—bad luck that affects operations, not just the bank account—is zero already. If so, count your blessings: hedging will probably not add any value. But many firms are not in so comfortable a position. To them hedging adds value because it improves the future cash-flow prospects from other activities. We return to this in Chapter 12.

4.5 CFO's Summary

In this chapter, we have analyzed forward contracts in a perfect market. We have discovered that forward contracts are essentially packaged deals, that is, transactions that are equivalent to a combination of a loan in one currency, a spot transaction, and a deposit in the other currency. In this sense, the forward contract is a distant forerunner of financial engineering. We have also seen how exchange markets and money markets are interlinked and can be used for arbitrage transactions and for identifying & comparing the two ways to make a particular transaction.

In perfect markets, it does not matter whether one uses forward contracts as opposed to their money-market replications. This holds for any possible transaction and its replication. For instance, a German firm will neither win nor lose if it replaces a EUR deposit by a swapped USD deposit since, from Interest Rate Parity, the two are equivalent. Or, more precisely, *if* it matters, it's because of market imperfections like spreads, or because the firm's other cash flows are affected too, but not because of the pure exchange of a FC cash flow by one in HC. We turn to market imperfections in the next chapter, and to the relevance of hedging in Chapter 12.

We have also found that the value of a forward contract is zero. This means that, everything else being the same, our German firm will not win or lose if it replaces a EUR deposit by an *uncovered* USD deposit. Again, a big word of caution is in order, here, because the "everything else being the same" clause is crucial. The above statement is perfectly true about the pure PV of two isolated cash flows, one in EUR and one in USD. But if the firm is so levered, the USD deposit is so large, and the EUR/USD so volatile that the investment could send the firm into receivership, then the dollar deposit would still not be a good idea—not because of the deposit *per se*, but because of the repercussions it could have on the firm's legal fees and interest costs and asset values. In short, the deposit's cash flows can have interactions with the company's other business, and these interactions might affect the firm's value.

A last crucial insight is that the forward rate is the market certainty equivalent, that is, the market's expectation corrected for any risks it thinks to be relevant. This insight can save a company a lot of time. It is also fundamental for the purpose of asset pricing. For cashflows with a known FC component, the logic is of course straightforward: (a) an asset with a known FC flow C_T^* is easy to hedge: sell forward C_T^* units of FC; (b) the hedged asset is easy to value; $(C_T^* \times F_{t,T})/(1 + r_{t,T})$; and (c) the unhedged asset must have the same value because the hedge itself has zero initial value and because a risk-free FC amount C_T^* cannot be affected by the hedge. Interestingly, under some distributional assumptions we can also apply the logic to cashflows that are highly non-linear functions of the future exchange rate. We return to this issue in Chapter 9.

Forward currency contracts have been around for centuries. A more recent instrument is the forward or futures contract on interest rates. Since forward interest contracts are not intrinsically 'international' and many readers may already know them from other sources, I stuff them into appendices, but if they new to you be warned that we are going to use them further on in this book. A key insight is that interest rates (spot and forward interest rates, and "yields at par") are all linked by arbitrage. Forward interest rates in various currencies are likewise linked through the forward markets.

4.6 Appendix: Interest Rates, Returns, and Bond Yields

4.6.1 Links Between Interest Rates and Effective Returns

We have defined the effective (rate of) return as the percentage difference between the initial (time-t) value and the maturity (time-T) value of a nominally risk-free asset over a certain holding period. For instance, suppose you deposit CLP 100,000 for six months, and the deposit is worth CLP 105,000 at maturity. The six-month effective return is:

$$r_{t,T} = \frac{105,000 - 100,000}{100,000} = 0.05 = 5\%. \tag{4.30}$$

In reality, bankers never quote effective rates of returns; they quote interest rates. An interest rate is an annualized return, that is, a return extrapolated to a twelve-month horizon. In the text, we emphasize this by adding an explicit *per annum* (or *p.a.*) qualification whenever we mention an interest rate. However, annualization can be done in many ways. It is also true that, for any system, there is a corresponding way to de-annualize the interest rate into the effective return—the number you need.

1. Annualization can be "simple" (*i.e.* linear): 5 percent for six months is extrapolated linearly, to 10 percent *per annum* (*p.a.*). A simple interest rate is the standard method for term deposits and straight loans when the time to maturity is less than one year. Conversely, the effective return is computed from the quoted simple interest rate as

$$1 + r_{t,T} = 1 + (T - t) \times [\text{simple interest rate}]. \tag{4.31}$$

Example 4.21

Let $(T - t)$ =1/2 year, and the simple interest rate 10 percent *p.a.* Then:

$$1 + r_{t,T} = 1 + 1/2 \times 0.10 = 1.05. \tag{4.32}$$

2. Annualization can also be compounded, with a hypothetical reinvestment of the interest. Using this convention, an increase from 100 to 105 in six months would lead to a constant-growth-extrapolated value of $105 \times 1.05 = 110.25$ after another six months. Thus, under this convention, 5 percent over six months corresponds to 10.25 percent *p.a.*. Conversely, the return is computed from the quoted compound interest rate as:

$$1 + r_{t,T} = (1 + [\text{compound interest rate}])^{T-t}. \tag{4.33}$$

Example 4.22

Let $(T-t)$ =1/2, and the compound interest rate 10.25 percent *p.a.*; then

$$1 + r_{t,T} = 1.1025^{1/2} = \sqrt{1.1025} = 1.05. \tag{4.34}$$

Compound interest is the standard method for zero-coupon loans and investments (without interim interest payments) exceeding one year.

3. Banks may also compound the interest every quarter, every month, or even every day. The result is an odd mixture of linear and exponential methods. If the interest rate for a six-month investment is i *p.a.*, compounded m times per year, the bank awards you i/m per subperiod of $1/m$ year. For instance, the *p.a.* interest rate may be $i = 6$ percent, compounded four times per year. This means you get $6/4 = 1.5$ percent per quarter. Your investment has a maturity of six months, which corresponds to two capitalization periods of one quarter each. After compounding over these two quarters, an initial investment of 100 grows to $100 \times (1.015)^2 = 103.0225$, implying an effective rate of return of 3.0225 percent. Thus, the effective return is computed from the quoted interest rate as:

$$1 + r_{t,T} = \left(1 + \frac{[\text{quoted interest rate}]}{m}\right)^{(T-t)m}. \tag{4.35}$$

Example 4.23

Let $(T-t)$ =1/2, and the compound interest rate 9.878% with quarterly compounding; then

$$1 + r_{t,T} = (1 + 0.098781/4)^{1/2 \times 4} = 1.05. \tag{4.36}$$

You may wonder why this byzantine mixture of linear and exponential is used at all. In the real world it is used when the bank has a good reason to understate the effective interest rate. This is generally the case for loans. For deposits, the reason may be that the quoted rate is capped (by law, like the U.S.' former Regulations Q and M; or because of a cartel agreement amongst banks). In finance theory, the mixture of linear and exponential is popular in its limit form, the continuously compounded rate:

4. In the theoretical literature, the frequency of compounding is often carried to the limit ("continuous compounding", *i.e.* $m \to \infty$). From your basic math course, you may remember that:

$$\lim_{m\to\infty} (1 + x/m)^m = e^x, \tag{4.37}$$

where e = 2.7182818 is the base of the natural (Neperian) logarithm. Conversely, the return is computed from the quoted interest rate ρ as:

$$1 + r_{t,T} = \lim_{m\to\infty}\left(1 + \frac{\rho}{m}\right)^{(T-t)m} = e^{\rho\,(T-t)} \tag{4.38}$$

Example 4.24

Let $(T - t) = 1/2$, and assume the continuously compounded interest rate equals 9.75803 percent. Then:

$$1 + r_{t,T} = e^{0.0975803/2} = 1.05. \tag{4.39}$$

Note the following link between the continuously and the annually compounded rates i and ρ:

$$(1 + i) = e^{ln(1+i)} \Rightarrow (1 + i)^{T-t} = e^{ln(1+i)\cdot(T-t)} \Rightarrow ln(1 + i) = \rho. \tag{4.40}$$

5. Bankers' discount is yet another way of annualizing a return. This is often used when the present value is to be computed for T-bills, promissory notes, and so on—instruments where the time-T value (or "face value") is the known variable, not thePV like in the case of a deposit or a loan. Suppose the time-T value is 100, the time to maturity is 0.5 years, and the p.a. discount rate is 5 percent. The present value will then be computed as

$$PV = 100 \times (1 - 0.05/2) = 97.5. \tag{4.41}$$

Conversely, the return is found from the quoted bankers' discount rate as:

$$1 + r_{t,T} = \frac{1}{1 - (T - t) \times [\text{banker's discount rate}]}. \tag{4.42}$$

Example 4.25

Let $(T - t) = 1/2$ and the *p.a.* bankers' discount rate 9.5238 percent. Then:

$$1 + r_{t,T} = \frac{1}{1 - 1/2 \times 0.095238} = 1.05. \tag{4.43}$$

In summary, there are many ways in which a bank can tell its customer that the effective return is, for instance, 5 percent. It should be obvious that what matters is the effective return, not the stated *p.a.* interest rate or the method used to annualize the effective return. For this reason, in most of this text, we use effective returns. This allows us to write simply $(1 + r_{t,T})$. If we had used annualized interest rates, all formulas would look somewhat more complicated, and would consist of many versions, one for each possible way of quoting a rate.

4.6.2 Common Pitfalls in Computing Effective Returns

To conclude this Appendix we describe the most common mistakes when computing effective returns. The first is forgetting to de-annualize the return. Always convert the bank's quoted interest rate into the effective return over the period $(T-t)$. *And* use the correct formula:

Example 4.26

Let $T-t = 3/4$ years. What are the effective rates of return when a banker quotes a 4 percent *p.a.* rate, to be understood as, alternatively, (1) simple interest, (2) standard compound interest, (3) interest compounded quarterly, (4) interest compounded monthly, (5) interest compounded daily, (6) interest compounded a million times a year, (7) interest compounded continuously, and (7) bankers' discount rate?

convention	formula	result (1+r)
simple	$1 + 3/4 \times 0.04$	1.030000000
compound, $m = 1$	$(1 + 0.04)^{3/4}$	1.029852445
compound, $m = 4$	$(1 + 0.04/4)^{4*3/4)}$	1.030301000
compound, $m = 12$	$(1 + 0.04/12)^{12*3/4}$	1.030403127
compound, $m = 360$	$(1 + 0.04/360)^{360*3/4}$	1.030452817
compound, $m = 1,000,000$	$(1 + 0.04/10^6)^{10^6*3/4}$	1.030454533
continuous compounding	$e^{0.04*3/4}$	1.030454533
banker's discount	$1/(1 - 3/4 * 0.04)$	1.030927835

Second, it is important to remember that there is an interest rate (or a discount rate) for every maturity, $(T-t)$. For instance, if you make a twelve-month deposit, the *p.a.* rate offered is likely to differ from the *p.a.* rate on a six-month deposit. Students sometimes forget this, because basic finance courses occasionally assume, for expository purposes, that the *p.a.* compound interest rate is the same for all maturities. Thus, there is a second pitfall to be avoided—using the wrong rate for a given maturity.

The third pitfall is confusing an interest rate with an internal rate of return on a complex investment. Recall that the return is the simple percentage difference between the maturity value and the initial value. This assumes that there is only one future cash flow. But many investments and loans carry numerous future cash flows, like quarterly interest payments and gradual amortisations of the principal. We shall discuss interest rates on multiple-payment instruments in the next appendix. For now, simply remember that the interest rate on, say, a five-year loan with annual interest payments should not be confused with the interest rate on a five-year instrument with no intermediate interest payments (zero-coupon bond).

Example 4.27

If a newspaper says the 10-year bond rate is 6%, this means that a bond with an annual coupon of 6% can be issued at par. That is, the 6 % is a "yield at par" on bullet bonds with annual coupons. What we need, in this chapter, are zero-coupon rates rather than yields at par.

4.7 Appendix: The Forward Forward and the Forward Rate Agreement

4.7.1 Forward Contracts on Interest Rates

You may know that loans often contain options on interest rates (caps and floors; see Chapter 16). Besides interest-rate options, there are also forward contracts on interest rates. Such forward contracts come under two guises: the Forward Forward contract (FF), and the Forward Rate Agreement (FRA).

A *Forward Forward contract* is just a forward deposit or loan: it fixes an interest rate today ($=t$) for a deposit or loan starting at a future time $T_1(>t)$ and expiring at $T_2(>T_1)$.

Example 4.28

Consider a six-to-nine-month Forward Forward contract for 10m Brazilian Real at 10 percent *p.a.* (simple interest). This contract guarantees that the return on a three-month deposit of BRL 10m, to be made six months from now, will be 10%/4 = 2.5%. At time T_1 (six months from now), the BRL 10m will be deposited, and the principal plus the agreed-upon interest of 2.5 percent will be received at time T_2 (nine months from now).

A more recent, and more popular, variant is the *Forward Rate Agreement.* Under an FRA, the deposit is *notional*—that is, the contract is about a hypothetical deposit rather than an actual deposit. Instead of effectively making the deposit, the holder of the contract will settle the gain or loss in cash, and pay or receive the present value of the difference between the contracted forward interest rate and market rate that is actually prevailing at time T_1.

Example 4.29

Consider a nine-to-twelve-month CAD 5m notional deposit at a forward interest rate of 4 percent *p.a.* (that is, a forward return of 1 percent effective). If the Interbank Offer Rate after nine months (T_1) turns out to be 3.6 percent *p.a.* (implying a return of 0.9 percent), the FRA has a positive value equal to the difference between the promised interest (1 percent on CAD 5m) and the interest in the absence of the FRA, 0.9 percent on CAD 5m. Thus, the investor will receive the present value of this contract, which amounts to:

$$\text{market value FRA} = \frac{5m \times (0.01 - 0.009)}{1.009} = 4955.40. \tag{4.44}$$

In practice, the reference interest rate on which the cash settlement is based is computed as an average of many banks' quotes, two days before T_1. The contract stipulates how many banks will be called, from what list, and how the averaging is done. In the early eighties, FRAs were quoted for short-term maturities only. Currently, quotes extend up to ten years.

4.7.2 Why FRAs Exist

Like any forward contract, an FRA can be used either for hedging or for speculation purposes. Hedging may be desirable in order to facilitate budget projections in an enterprise or to reduce uncertainty and the associated costs of financial distress. Banks, for example, use FRAs, along with T-bill futures and bond futures, to reduce maturity mismatches between their assets and liabilities. For instance, a bank with average duration of three months on the liability side and twelve months on the asset side, can use a three-to-twelve month FRA to eliminate most of the interest risk. An FRA can, of course, serve as a speculative instrument too.

As we shall show in the next paragraph, FFs (or FRAs) can be replicated from term deposits and loans. For financial institutions, and even for other firms, FRAs and interest futures are preferred over such synthetic FRAs in the sense that they do not inflate the balance sheet.

Example 4.30

Suppose that you need a three-to-six month forward loan for JPY 1b. Replication would mean that you borrow (somewhat less than) JPY 1b for six months and invest the proceeds for three months, until you actually need the money. Thus, your balance sheet would have increased by JPY 1b, without any increase in profits or cash flows compared to the case where you used a Forward Forward or an FRA.

The drawback of using an FF or FRA is that there is no organized secondary market. However, as in the case of forward contracts on foreign currency, long-term FRA contracts are sometimes collateralized or periodically recontracted. This reduces credit risk. Thus, a fairly active over-the-counter market for FRAs is emerging.

4.7.3 The Valuation of FFs (or FRAs)

We now discuss the pricing of FFs (or FRAs—both have the same value): How should one value an outstanding contract, and how should the market set the normal forward interest rate at a given point in time? In this section, we adopt the following notation:

t_0 : the date on which the contract was initiated

$t\ (\geq t_0)$: the moment the contract is valued

T_1: the expiration date of the forward contract (that is, the date that the gains or losses on the FRA are settled, and the date at which the notional deposit starts)
$T_2\ (> T_1)$: the expiration date of the notional deposit

$r^f_{t_0,T_1,T_2}$: the effective return between T_1 and T_2, without annualization, promised on the notional deposit at the date the FRA was signed, t_0.

First consider a numerical example:

DoItYourself problem 4.6

Consider a FF under which you will deposit JPY 1b in nine months and receive 1.005b in twelve. The effective risk-free rates for these maturities are r_{t,T_1}=0.6% and r_{t,T_2}=0.81, respectively. Value each of the PN's that replicate the two legs of the FF. Compute the net value.

The generalisation is obvious. Below, we take a notional deposit amount of 1 (at T_1):

$$\begin{aligned}
\text{PV, at } t\text{, of a unit FF} &= \underbrace{\frac{\text{promised inflow at } T_2}{1+r_{t,T_2}}}_{\text{“asset” PN}} - \underbrace{\frac{\text{promised outflow at } T_1}{1+r_{t,T_1}}}_{\text{“liability” PN}} \\
&= \frac{1+r^f_{t_0,T_1,T_2}}{1+r_{t,T_2}} - \frac{1}{1+r_{t,T_1}}. \qquad (4.45)
\end{aligned}$$

In one special case we can consider the expiry moment ($t = T_1$):

DoItYourself problem 4.7

Derive, from this general formula, our earlier cash-settlement equation,

$$\text{PV, at } T_1\text{, of a unit FF} = \frac{r^f_{t_0,T_1,T_2} - r_{T_1,T_2}}{1+r_{T_1,T_2}}. \qquad (4.46)$$

The other special case worth considering is the value at initiation ($t_0 = t$). We know that this value must be zero, like for any standard forward contract, so this provides a way to relate the forward rate to the two spot rates, all at t:

DoItYourself problem 4.8

Derive, from the general formula, the relation between the time-t spot and forward rates:

$$(1+r_{t,T_1})(1+r^f_{t,T_1,T_2}) = 1+r_{t,T_2}; \qquad (4.47)$$

$$\Leftrightarrow r^f_{t,T_1,T_2} = \frac{1+r_{t,T_2}}{1+r_{t,T_1}} - 1. \qquad (4.48)$$

The left-hand side of the first equality, Equation [4.47], has an obvious interpretation: it shows the gross return from a synthetic deposit started right now (t) and expiring at T_2, made not directly, but replicated by making a t-to-T_1 spot deposit which is rolled over (i.e. re-invested, here including the interest earned) via a T_1-to-T_2 forward deposit. So the money is contractually committed for the total t-to-T_2 period, and the total return is fixed right now—two ingredients that also characterize a t-to-T_2 deposit. In that light, Equation [4.47] just says that the direct and the synthetic t-to-T_2 deposits should have the same return.

Figure 4.6: **Spot and Forward Money Markets (with International Links)**

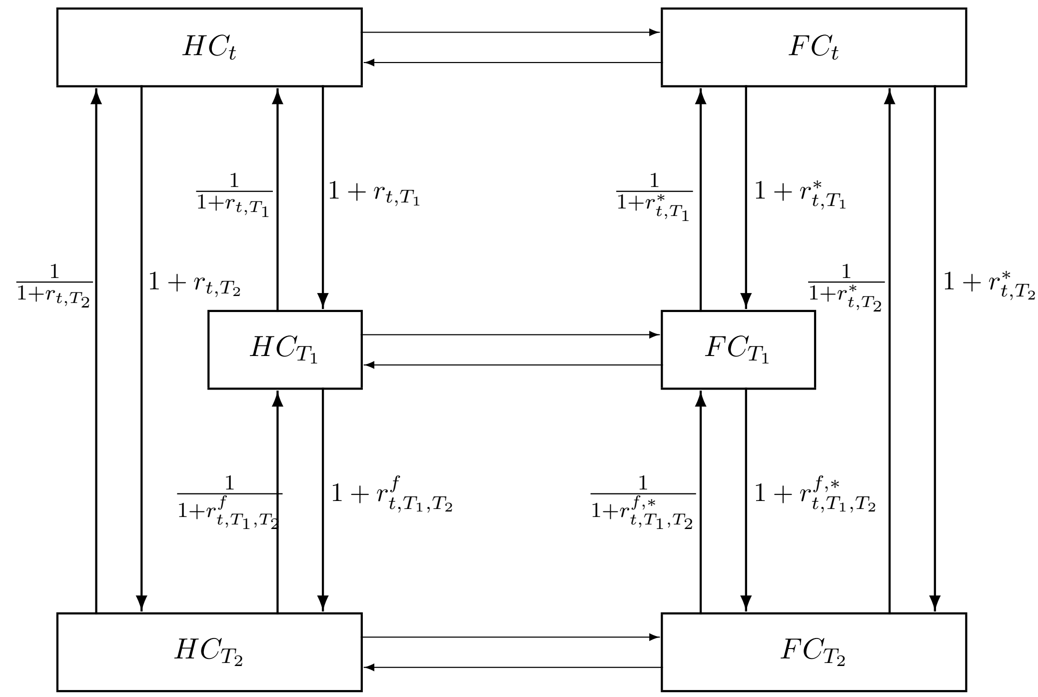

As in the case of currency forwards, no causality is implied by our way of expressing Equation [4.48]. The three rates are set jointly and have to satisfy Equation [4.48], that's all. As in Chapter 4, one could argue that causality, if any, may run from the forward interest rate towards the spot rate because the forward rate reflects the risk-adjusted expectations about the future interest rate. We shall use Equation [4.48] when we discuss eurocurrency futures, in the Appendix to Chapter 6.

There is an obvious no-arbitrage version of this. In Figure 4.6 we combine two of our familiar spot-forward currency diagrams, one for future date T_1 and the other for date T_2. The focus, this time, is not on the exchange markets, so the horizontal lines that refer to currency deals are made thinner. The forward deposits and loans are shown as transactions that transform T_1-dated money into T_2 money (the deposit) or vice versa (the loan), and the multiplication factors needed to compute the output from a transaction, shown next to the arrows, are $(1+r^f)$ and $1/(1+r^f)$, respectively. This diagram shows that every spot or forward money-market deal can be replicated, which helps you in shopping-around problems. The diagram also helps identifying the no-arb constraints.

DoItYourself problem 4.9

We have already shown how to replicate the t-to-T_2 deposit. In the table below, add the replications for the other transactions and check that they generate the gross returns shown in the rightmost column.

replicand	replication	output value#
t-to-T_2 deposit	spot deposit t to T_1, rolled over forward T_1 to T_2	$(1+r_{t,T_1})(1+r^f_{t,T_1,T_2})$
forward deposit T_1 to T_2		$\frac{1+r_{t,T_2}}{1+r_{t,T_1}}$
spot deposit t to T_1		$\frac{1+r_{t,T_2}}{1+r^f_{t,T_1,T_2}}$
t-to-T_2 loan		$\frac{1}{(1+r_{t,T_1})(1+r^f_{t,T_1,T_2})}$
forward loan T_1 to T_2		$\frac{1+r_{t,T_1}}{1+r_{t,T_2}}$
spot loan t to T_1		$\frac{1+r^f_{t,T_1,T_2}}{1+r_{t,T_2}}$

#: output is computed for input value equal to unity.

.

But the diagram shows not just the replication possibilities: there are also two no-arb constraints inside each money market, corresponding to, respectively, the clockwise and counterclockwise roundtrips. You start for instance in the time-T_2 box, issuing a PN note dated T_2. You discount it immediately and make a synthetic deposit t to T_2. The constraint is that the proceeds of this deposit be no higher than 1, the amount you owe to the holder of the PN:

$$\begin{gathered} HC_{T_2} \to HC_t \to HC_{T_1} \to HC_{T_2} \\ 1 \times \frac{1}{1+r_{t,T_2}} \times (1+r_{t,T_1}) \times (1+r^f_{t,T_1,T_2}) \leq 1. \end{gathered} \tag{4.49}$$

DoItYourself problem 4.10

Identify the other no-arb trip in the home money market and write the corresponding constraint. Combine it with constraint [4.49] and check that you find back Equation [4.47].

To seasoned arbitrageurs like you it is easy to add bid- and ask-superscripts to the rates of return (and to the exchange rates, while you are at it). The no-arb constraints still are [synthetic bid] ≤ [ask] and [bid] ≤ [synthetic ask], with the synthetics computed from the worst-possible-combination versions of the perfect-market replication that you just worked out yourself.

Before we move to other markets, there is another set of no-arb constraints and shopping-around opportunities to be discussed, namely those created via international linkages rather than relations within each money market. Remember one can replicate a currency-X spot deposit or loan by swapping a currency-Y spot deposit or loan into currency X. Well, the same holds for forward deposits and loans. For instance, in the few years when USD or GBP had FRA markets but minor European currencies had not (yet), pros replicated the missing FRA's by swapping USD or GBP FRA's into, say, NLG via a forward-forward currency swap, in or out. Such swaps are described in Chapter 5, and consist of a currency forward in one direction combined with a second currency forward, in the other direction. In short, when the starting date of a deposit or loan is not spot but n days forward, we just replace the spot leg of the swap by the appropriate forward leg.

4.7.4 Forward Interest Rates as the Core of the Term Structure(s)

Remember that forward exchange rates, being the risk-adjusted expectations, are central in any theory of exchange rates. In the same way, forward interest rates can be viewed as the core of every theory of interest rates. The standard *expectations theory* hypothesizes that forward interest rates are equal to expected future spot rates, and Hicks added a *risk premium*, arguing—to use a post-Hicksian terminology—that the beta risk of a bond is higher the longer its time to maturity. Modern versions would rather state everything in terms of PN prices rather than interest rates, but would agree with the basic intuition of the old theories: forward rates reflect expectations corrected for risks.

Various theories or models differ as to how expectations evolve and risk premia are set, but once the forward rates are set, the entire term structure follows. We illustrate this with a numerical example, and meanwhile initiate you to the various interest-rate concepts: spot rates, yields at par for bullet bonds, and other yields at par.

We start from the first row in Table 4.3, which shows a set of forward rates. For simplicity of notation, current time t is taken to be zero, so that a one-period forward rate looks like $r^f_{0,n-1,n}$ rather than the more laborious $r^f_{t,t+n-1,t+n}$. For some reason—mainly expectations, one would presume—there is a strong "hump" in the forward rates: they peak at the 3-to-4 year horizon. (A period is of unspecified length, in the theories; but let's agree they are years).[10] The initial spot rate and the forward rate with starting date 0 are, of course, the same. Below we show you the formulas to be used in a spreadsheet to generate all possible term structures (TS).

The TS of spot rates is obtained in two steps. First we cumulate the forward

[10]For this reason the only non-arbitrary theory is one that works with continuous time, where a period lasts dt years. But for intro courses this has obvious drawbacks.

Table 4.3: **Term Structures and their Linkages**

	rate or PV factor for $n = ...$						
	1	2	3	4	5	6	7
forward rate *p.p.*, $r^f_{0,n-1,n}$	0.0300	0.0350	0.0380	0.0400	0.0360	0.0300	0.0200
$1 + r^f_{0,n-1,n}$	1.0300	1.0350	1.0380	1.0400	1.0360	1.0300	1.0200
$1 + r_{0,n} = \Pi^n_{j=1}(1 + r^f_{0,j-1,j})$	1.0300	1.0661	1.1066	1.1508	1.1923	1.2280	1.2526
$\bar{r}_{0,n} = (1 + r_{0,n})^{1/n} - 1$	0.0300	0.0325	0.0343	0.0357	0.0358	0.0348	0.0327
$\mathrm{PV}_{0,n} = 1/(1 + r_{0,n})$	0.9709	0.9380	0.9037	0.8689	0.8387	0.8143	0.7984
PV annuity, $a_{0,n} = \sum^n_{j=1} PV_{0,j}$	0.9709	1.9089	2.8126	3.6816	4.5203	5.3346	6.1330
$R_{0,n} : \frac{[1-(1+R_{0,n})^{-n})}{R_{0,n}} = a_{0,n}$	0.0300	0.0316	0.0330	0.0340	0.0346	0.0347	0.0342
$c_{0,n} = (1 - \mathrm{PV}_{0,n})/a_{0,n}$	0.0300	0.0325	0.0342	0.0356	0.0357	0.0348	0.0329

Key Starting from an assumed set of forward rates I compute the set of 'spot' zero-coupon rates (lines 3 and 4) and present value factors (line 5). This allows us to find the PV of a constant unit annuity (line 6) and the corresponding yield. Finally I compute the yield at par for a bullet bond. The math is described in the text.

rates into effective spot rates, using Equation [4.47]:

$$1 + r_{0,n} = \Pi^n_{j=1}(1 + r^f_{0,j-1,j}). \tag{4.50}$$

The rate on the left-hand side is the effective rate we have always used in this book. But for the theory of term structures it is useful to convert the effective rate to a per-period rate, which we denote by $\bar{r}$. The computation is

$$1 + \bar{r}_{0,n} := (1 + r_{0,n})^{1/n}. \tag{4.51}$$

The spot rates are the yields to maturity on zero-coupon bonds expiring at n. Note how the per-period gross rates are rolling geometric averages—numerically close to simple averages—of all gross forward rates between times 0 and n.[11] See how the strong hump is very much flattened out by the rolling-averaging, and the peak pushed to $n = 5$ instead of $n = 4$ for the forward rates. A second alternative way to work with the effective rate is to compute the PV of one unit of HC payable at time n,

$$\mathrm{PV}_{0,n} = 1/(1 + r_{0,n}). \tag{4.52}$$

The TS of yields for constant-annuity cash flows is a different TS. It is not as popular as the TS of yields at par for bullet loans, see below, but it is convenient to look at this one first. Any yield or *internal rate of return* is the compound "flat" rate that equates a discounted stream of known future cash flows C_j to an observed present value:

$$y : \frac{C_1}{1+y} + \frac{C_2}{(1+y)^2} + ... + \frac{C_n}{(1+y)^n} = \text{observed PV}. \tag{4.53}$$

[11] A gross rate is $1 + r$, r being the net rate we always use in this text.

Figure 4.7: **Term Structures: Forward, Spot, and Two Types of Yields**

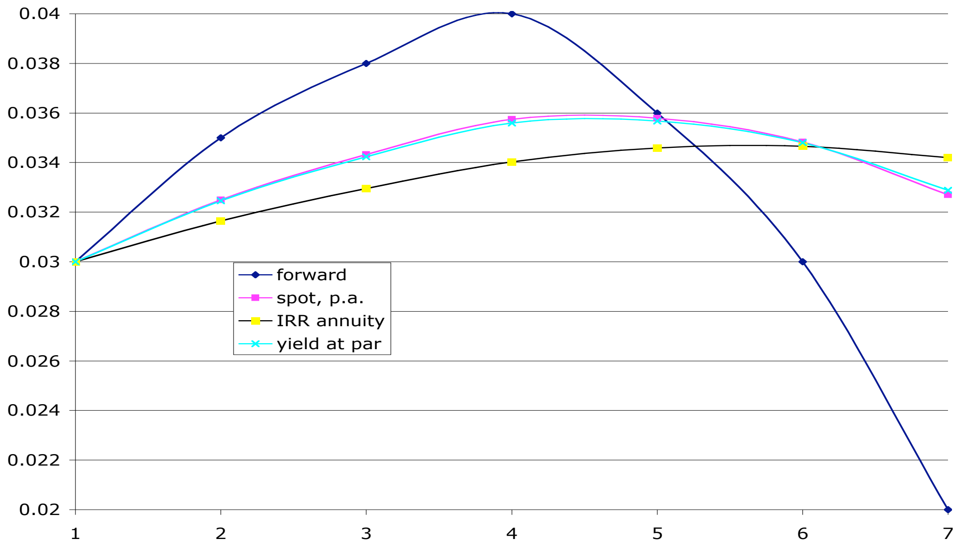

Here we look at the special case $C_j = 1, \forall j$, the constant unit-cash-flow stream, the right-hand side of the above equation. Let us first find thePV of the constant stream. Since we already know the PV of a single unit payment made at n, the PV of a stream paid out at times 1, ... , n is simply the sum. This special PV is denoted as $a_{0,n}$, from "annuity". We compute its value for various n as

$$a_{0,n} = \sum_{j=1}^{n} PV_{0,j}. \tag{4.54}$$

Next we find the yield that equates this PV to the discounted cash flows. When the cash flows all equal unity, the left-hand side of Equation [4.53] is equal to $(1-(1+y)^{-n})/y$, but the y that solves the constraint must still be found numerically, using e.g. a spreadsheet tool. In the table the result is found under the label $R_{0,n}$. Note how this yield is an analytically non-traceable mixture of all spot rates. The hump is flattened out even more, and its peak pushed back one more period.

The TS of yields at par for bullet loans is defined as a yield that sets the PV of a bullet loan equal to par. But it is known that to get a unit value the yield must be set equal to the coupon rate. So we can now rephrase the question as follows: how do we set the coupon rate c such that the PV's of the coupons and the principal sum to unity?

$$c_{0,n} : \underbrace{c_{0,n} \times a_{0,n}}_{\text{PV of coupons}} + \underbrace{\text{PV}_{0,n} \times 1}_{\text{PV of amortization}} = 1 \Rightarrow c_{0,n} = \frac{1 - \text{PV}_{0,n}}{a_{0,n}}. \tag{4.55}$$

Again, this is numerically much closer to the spot rates than the yield on constant-annuity loans, and the reason obviously is that the bullet loan is closer to a zero-

Figure 4.8: **Extracting spot and forward rates from the JPY swap rates**

	swap % p.a	spot % p.a	forwd % p.a
1	0.610	0.610	0.610
2	0.935	0.937	1.264
3	1.195	1.199	1.727
4	1.405	1.413	2.058
5	1.580	1.593	2.316
6	1.725	1.744	2.499
7	1.855	1.880	2.701
8	1.965	1.996	2.814
9	2.050	2.086	2.812
10	2.125	2.167	2.894

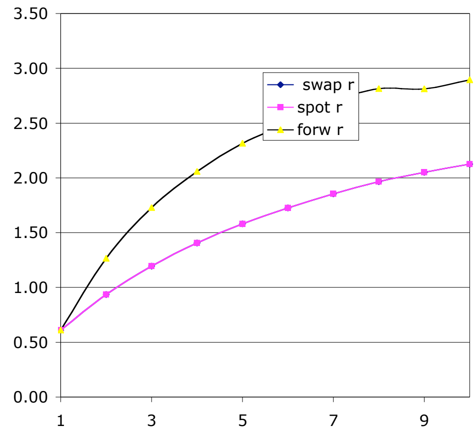

Key Swap and spot are so close, in these figures, that on the graph you no longer spot the difference; you need to look at the numbers.

coupon loan—especially in an example where, like in ours, interest rates are generally low. In Figure 4.7, which shows the four term structures graphically, those for swap and spot rates overlap almost perfectly.

This illustrates how the TS of forward rates contains all information for pricing, so that TS theories are basically theories about forward rates. It also gives you a feeling how swap dealers set their long-term interest rates, yields-at-par for bullet bonds. Like us here, they construct them from spot rates. These spot rates, in turn, are obtained from PV factors extracted, via regression analysis, from bond prices in the secondary market. You can, of course, reverse-engineer all this and extract PV factors from swap rates, and thence forward rates. Then you may ask the question whether there seem to be good reasons for the forward rate to behave as it appears to do, and perhaps invest or disinvest accordingly. For instance, in Figure 4.8 we have taken the JPY swap rates from Chapter 7 and extracted spot and forward rates. Spot rates are familiarly close to swap rates (yields at par for bullet bonds) but the forward rate, equally familiar, moves much faster than the spot rate (a rolling average). So one can ask the question how these forward rates compare to your expectations about future spot rates.

A second insight you should remember is that there is no such thing as "the" TS. Academics would first think of the TS of spot rates or forward rates (and be precise about that). But practitioners first think about the TS of yields at par for bullet bonds, the numbers one sees in the newspaper or that are quoted by swap dealers (who call them swap rates). Many traditional practitioners would apply the yield-at-par rates for any instrument, whether it is a bullet loan or not. That can imply serious errors and inconsistencies.

Yields are funny. Even if we just consider bullet bonds, there still is a yield for every *total* time to maturity n. So in the situation depicted in Table 4.3 a coupon paid at time 1 would be discounted at 3% if it is part of a one-period bond, 3.25% if it is part of a two-period bond, 3.42% if it is part of a three-period bond, and so on. It is much more logical to work with a discount rate for every payment horizon, regardless of what bond pays out the money, rather than a discount rate for every bond, regardless of the date of the payment.

4.8 Test Your Understanding

4.8.1 Quiz Questions

1. Which of the following statements are correct?

   (a) A forward purchase contract can be replicated by: borrowing foreign currency, converting it to domestic currency, and investing the domestic currency.

   (b) A forward purchase contract can be replicated by: borrowing domestic currency, converting it to foreign currency, and investing the foreign currency.

   (c) A forward sale contract can be replicated by: borrowing foreign currency, converting it to domestic currency, and investing the domestic currency.

   (d) A forward sale contract can be replicated by: borrowing domestic currency, converting it to foreign currency, and investing the foreign currency.

   (e) In a perfect market you could forbid forward markets (on the basis of anti-gambling laws, for instance), and nobody would give a fig.

   (f) The spot rate and the interest rate determine the forward price.

   (g) No, the forward determines the spot.

   (h) No, the forward and the spot and the foreign interest rate determine the domestic interest rate.

   (i) No, there are just four products that are so closely related that their prices cannot be set independently.

2. What's wrong with the following statements?

   (a) The forward is the expected future spot rate.

   (b) The sign of the forward premium tells you nothing about the strength of a currency; it just reflects the difference of the interest rates.

   (c) The difference of the interest rates tells you nothing about the strength of a currency; it just reflects the forward premium or discount.

   (d) The forward rate is a risk-adjusted expectation but the spot rate is independent of expectations

   (e) A certainty equivalent tends to be above the risk-adjusted expectation because of the risk correction.

   (f) A risk-adjusted expectation is always below the true expectation because we don't like risk.

   (g) A risk-adjusted expectation can be close to, or above the true expectation. In that case the whole world would hold very little of that currency, or would even short it.

   (h) Adding a zero-value contract cannot change the value of the firm; therefore a forward hedge cannot make the shareholders better off.

4.8.2 Applications

1. Check analytically the equivalence of the two alternative ways to do the following trips:

   (a) Financing of international trade: you currently hold a FC claim on a customer payable at T, but you want cash HC instead.

   (b) Domestic deposits: you currently hold spot HC and you want to park that money in HC, risk-free.

   (c) You want to borrow HC for 3 months.

   (d) Immunizing a HC dent: you want to set aside some of your cash HC so as to take care of a future FC debt.

   (e) Borrowing FC. You want to borrow FC but a friend tells you that swapping a HC loan is much cheaper

2. You hold a set of forward contracts on EUR, against USD (=HC). Below I show you the forward prices in the contract; the current forward prices (if available) or at least the current spot rate and interest rates (if no forward is available for this time to maturity). Compute the fair value of the contracts.

   (a) Purchased: EUR 1m 60 days (remaining). Historic rate: 1.350; current rate for same date: 1.500; risk-free rates (simple per annum): 3% in USD, 4% in EUR.

   (b) Purchased: EUR 2.5m 75 days (remaining). Historic rate: 1.300; current spot rate: 1.5025; risk-free rates (simple per annum): 3% in USD, 4% in EUR.

   (c) Sold: EUR 0.75m 180 days (remaining). Historic rate: 1.400; current rate for same date: 1.495; risk-free rates (simple per annum): 3% in USD, 4% in EUR.

3. 60-day interest rate (simple, p.a.) are 3% at home (USD) and 4% abroad (EUR). The spot rate moves from 1.000 to 1.001.

   (a) What is the return differential, and what is the corresponding prediction of the change in the forward rate?

   (b) What is the actual change in the forward rate?

   (c) What is the predicted change in the swap rate computed from the return differential?

   (d) What is the actual change in the swap rate?

4. 60-day interest rate (simple, p.a.) are 3% at home (USD) and 4% abroad (EUR). The spot rate is 1.250.

   (a) Check that investing EUR 1m, hedged, returns as much as USD 1.25m

(b) Check that if taxes are neutral, and the tax rate is 30%, also the after-tax returns are equal. (Yes, this *is* trivial.)

(c) How much of the income from swapped EUR is legally interest income and how much is capital gain or loss?

(d) If you do not have to pay taxes on capital gains and cannot deduct capital losses, would you still be indifferent between USD deposits and swapped EUR?

5. 60-day interest rate (simple, p.a.) are 3% at home (USD) and 4% abroad (EUR). The spot rate is 1.250.

   (a) Check that borrowing EUR 1m (=current proceeds, not future debt), hedged, costs as much as borrowing USD 1.25m

   (b) Check that if taxes are neutral, and the tax rate is 30%, also the after-tax costs are equal. (Yes, this *is* trivial.)

   (c) How much of the costs of bowwowing swapped EUR is legally interest paid and how much is capital gain or loss?

   (d) If you do not have to pay taxes on capital gains and cannot deduct capital losses, would you still be indifferent between USD loans and swapped EUR?

6. Groucho Marx, as Governor of Freedonia's central bank, has problems. He sees the value of his currency, the FDK, under constant attack from Rosor, a wealthy mutual-fund manager. Apparently, Rosor believes that the FDK will soon devalue from GBP 1.000 to 0.950.

   (a) Currently, both GBP and FDK interest rates are 6% *p.a.* By how much should Groucho change the one-year interest rate so as to stabilize the spot rate even if Rosor expects a spot rate of 0.950 in one year? Ignore the risk premium—that is, take 0.950 to be the certainty equivalent.

   (b) If the interest-rate hike also affects Rosor's expectations about the future spot rate, in which direction would this be? Taking into account also this second-round effect, would Groucho have to increase the rate by more than your first calculation, or by less?

Chapter 5

Using Forwards for International Financial Management

In this chapter, we discuss the five main purposes for which forward contracts are used: arbitrage (or potential arbitrage), hedging, speculation, shopping around, and valuation. These provide the topics of Sections 5.2 to 5.6, respectively. But first we need to spend some time on practical issues: the quotation method, and the provisions for default risk (Section 5.1).

5.1 Practical Aspects of Forwards in Real-world Markets

5.1.1 Quoting Forward Rates with Bid-Ask Spreads

With bid-ask spreads, a forward rate can still be quoted "outright" (that is, as an absolute number), or as a swap rate. The outright quotes look like spot quotes in that they immediately give us the level of the forward bid and ask rates; for instance, the rates may be CAD/USD (180 days) 1.1875–1.1895. Swap rates, on the other hand, show the numbers that are to be added to/subtracted from the spot bid and ask rates in order to obtain the forward quotes. One ought to be careful in interpreting such quotes, and make sure that the correct number is added to or subtracted from the spot bid or ask rate.

Example 5.1

Most papers nowadays show outright rates, but Antwerp's *De Tijd* used to publish swap rates until late 2005. Table 5.1 shows an example, to which I added a column of midpoint swap rates and LIBOR 30d interest rates (simple, p.a.). Swap rates are

Figure 5.1: **Swap Quotes, bid and ask, from *De Tijd***

Termijnkoersen	1 maand		2 maand		3 maand		6 maand		12 maand	Bron: Dexia	Spot rate	LIBOR 30d
Amerikaanse dollar	19.20	19.28	37.30	37.50	58.82	59.07	115.00	115.60	229.60	231.00	1.1776	4.20
Australische dollar	46.00	46.60	84.20	85.10	128.00	130.00	239.00	242.00	464.00	468.00	1.5988	5.55
Brits pond	13.40	13.60	24.20	24.50	36.50	36.80	65.70	66.40	121.00	123.00	0.6846	4.81
Japanse Yen	-29.10	-28.80	-57.20	-56.80	-89.10	-88.60	-177.00	-176.00	-370.00	-366.00	139.7800	0.03
Nieuw-Zeelandse dollar	80.10	81.10	148.00	149.00	226.00	228.00	425.00	429.00	818.00	829.00	1.7035	7.48
Zweedse Kroon	-52.10	-47.80	-132.00	-126.00	-189.00	-181.00	-372.00	-356.00	-655.00	-607.00	9.5162	1.60
Zwitserse frank	-21.40	-21.10	-39.70	-38.90	-60.00	-59.70	-114.00	-111.00	-211.00	-205.00	1.5491	0.80
eur												2.335

quoted in foreign currency since the quotes against the euro are conventionally in FC units; and they are in basis points, *i.e.* hundredths of cents.

To compute the outright forward rates from these quotes, one adds the first swap rate to the spot bid rate, and the second swap rate to the spot ask rate. The excerpt shows the midpoint spot rate rather than the bid-ask quotes. Suppose, however, that the bid and ask spot rates are (1.17)74-78 for the USD. Then the outright forward rates, one month, are computed as follows:

– Bid: $1.1774 + 0.0001920 =$ USD/EUR 1.1775920.

– Ask: $1.1778 + 0.0001928 =$ USD/EUR 1.1779928.

DoItYourself problem 5.1

Check the interest rates, and note which ones are higher than the EUR one. Figure out which forward rates should be above par and which below. Verify that the signs of the swap rates are correct, especially once you remember that the EUR is the FC (also for the GBP quote).

Note from the example that in case of a premium we always add the smaller of the two swap rates to the spot bid rate, and the larger swap rate to the spot ask rate. As a result, the forward spread is wider than the spot spread (Figure 5.2). Likewise, in case of a discount, the number we subtract from the spot bid rate is larger, in absolute value, than the number we subtract from the spot ask rate; and this again produces a wider spread in the forward market than in the spot market (Figure 5.2). Finally, note that the difference between the swap rates becomes larger the longer the contract's time to maturity. This illustrates the *Second Law of Imperfect Exchange Markets*: the forward spread is always larger than the spot spread, and increases with the time to maturity.

One explanation of this empirical regularity is that the longer the maturity, the lower the transaction volume; and in thin markets, spreads tend to be high. A

Figure 5.2: **The bid-ask spread in a forward is wider than in a spot**

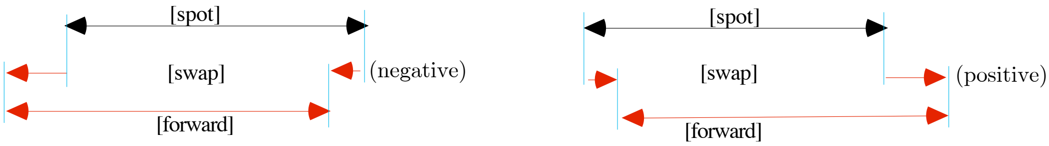

Key For negative swap rates the bid is the bigger one , in absolute terms, while for positive swap rates the ask is the bigger one. This is equivalent to observing a larger total bid-asp spread in the forward market.

second reason is that, over short periods, things generally do not change much, but a lot can happen over long periods. Thus, a bank may be confident that the customer will still be sound in 30 days, but feel far less certain about the customer's creditworthiness in five years. In addition, also the exchange rate can change far more over five years than over 30 days; so the farther the maturity date, the larger also the potential loss if and when default would happen and the bank would have to close out, *i.e.* reverse, the forward contract it had signed with the customer at t_0.[1] Thus, banks build in a default-risk premium into their spreads, which, therefore, goes up with time to maturity. Later on we will see by how much the spreads can go up maximally with time to maturity.

The Second Law keeps you from getting irretrievably lost when confronted with bid-ask swap quotes, because the convention of quoting is by no means uniform internationally. Sometimes the sign of the swap rate (+ versus −; or p versus d) is entirely omitted, because the pros all know that sign already. Or sometimes the swap rates are quoted, regardless of sign, as "small number–big number," followed by p (for premium) or d (for discount). When in doubt, just try which combination generates the bigger spread.

Let us now address weightier matters: how is credit risk handled?

5.1.2 Provisions for Default

Forward dealers happily quote forward rates based on interbank interest rates, even if their counterpart is much more risky than a bank. Shouldn't they build risk spreads into the interest rates, like when they lend money? The answer is No (or, at most, Hardly): while the bank's risk under a forward contract is not entirely absent, it is still far lower than under a loan contract. Banks have, in effect, come up with various solutions that partially solve the problem of default risk:

[1]Note that the exchange risk is only relevant if and when the customer defaults. Normally, a bank closes its position soon after the initial deal is signed, but this close-out position unexpectedly turns out to be an open one if and when the customer's promised deal evaporates. In short, exchange risk only arises as an interaction with default risk.

- **The right of offset** First and foremost, a forward contract has an unwritten but time-hallowed clause saying that if one party defaults, then the other party cannot be forced to do its own part of the deal; moreover, if that other party still sustains losses, the defaulting party remains liable for these losses. Thus, if the customer defaults, the bank that sold FC forward can now dispose of this amount in the spot market (rather than delivering it to the defaulting customer) and keep the revenue. There still is a potential loss if and to the extent that this revenue (S_T) would be below the amount promised ($F_{t_0,T}$), but even if nothing of this can be recouped in the bankruptcy court the maximum loss is ($F_{t_0,T} - S_T$), not $F_{t_0,T}$.[2]

  Example 5.2
  Citibank has sold forward JPY 100m at USD/JPY 0.0115 to Fab4 Inc, a rock band, to cover the expenses of their upcoming tour; but on the due date Citi discovers they have declared bankruptcy. Being a careful banker, City had bought forward the Yens it owed Fab4. Given the bankruptcy, Citi has no choice but to sell these JPY 100m spot at, say, S_T =0.0109. The default has cost Citi 100m × (0.0115-0.0109) = USD 60,000. In contrast, if Fab4 had promised JPY 100m in repayment of a loan, Citi might have lost the full 100m × 0.0115 = USD 1.15m. Since, under the forward contract, Citi can revoke its own obligation the net loss is always smaller, and could even turn into a gain.

- **Interbank: credit agreements** In the interbank market, the players deal only with banks and corporations that are well-known to one another and have signed credit agreements for (spot and) forward trading—that is, agreements that they will freely buy and sell to each other. Even there, credit limits are set per bank to limit default risk.

- **Firms: credit agreements or security** Likewise, corporations can buy or sell forward if they are well-known customers with a credit agreement providing—within limits—for spot and forward trades, probably alongside other things like overdraft facilities and envelopes for discounting of bills or for letters of credit. The alternative is to ask for margin. For unknown or risky customers, the margin may be as high as 100 percent.

  Example 5.3
  Expecting a depreciation of the pound sterling, Burton Freedman wants to sell forward GBP 1m for six months. The 180-day forward rate is USD/GBP 1.5. The bank, worried about the contingency that the pound may actually go up, asks for 25 percent margin. This means that Mr. Freedman has to deposit 1m x USD/GBP

[2] To obtain a security with the same credit risk in the case of a synthetic forward contract, the bank would have to insist that the customer hold the deposit part of its synthetic contract in an escrow account, to be released only after the customer's loan is paid back. The forward contract is definitely the simpler way to achieve this security—which is one reason why an outright contract is more attractive than its synthetic version.

1.5 × 0.25 = USD 375,000 with the bank, which remains with the bank until he has paid for the GBP. The interest earned on the deposit is Mr. Freedman's. This way, the bank is covered against the combined contingency of the GBP rising by up to 25 percent *and* Mr. Freedman defaulting on the contract.

- **Restricted use** Even within an agreed credit line, "speculative" forward positions are frowned upon—unless a lot of margin is posted. Banks see forwards as hedging devices for their customers, not as speculative instruments.
- **Short lives** Maturities go up to 10 years, but in actual fact the life of most forward contracts is short: most contracts have maturities less than one year, and longer-term contracts are entered into only with customers that have excellent credit ratings. To hedge long-term exposures one then needs to roll over short-term forward contracts. For example, the corporation can engage in three consecutive one-year contracts if a single three-year contract is not available.

  Example 5.4

  At time 0, an Indian company wants to buy forward USD 1m for three years. Suppose that the bank gives it a three-year forward contract at $F_{0,3}$ = INR/USD 40. Suppose the bank's worst dreams become true: the spot rate goes down all the time, say, to 38, 36, and 34 at times 1, 2, and 3, respectively. If, at time 3, the company defaults, the bank is stuck with USD 1m worth INR 34m rather than the contracted value, 40m. Thus, the bank has a loss of $(F_{0,3} - S_3)$ = INR 6m.

  Suppose instead that at t = 0, the bank gave a one-year contract at the rate $F_{0,1} = 40.3$. After one year, the customer pays INR 40.3m for the currency, takes delivery of the USD 1m, and sells these (spot) at $S_1 = 38$. After verification of the company's current creditworthiness, the bank now gives it a new one-year contract at, say, $F_{1,2} = 37.2$. At time t = 2, the customer takes the second loss. If it is still creditworthy, the customer will get a third one-year forward contract at, say, $F_{2,3} = 35.9$. If there is default at time 3, the bank's loss on the third contract is just 1.9m rather than the 6m it would have lost with the three-year contract.

From the bank's point of view, the main advantage of the alternative of rolling over short-term contracts is that losses do not accumulate. The uncertainty, at time 0, about the spot rate one year out is far smaller than the uncertainty about the rate three years out. Thus, *ex ante* the worst possible loss on a three-year contract exceeds the worst possible loss on a one-year contract. In addition, the probability of default increases with the time horizon—in the course of three years, a lot more bad things can happen to a firm than in one year, *ex ante*—and also with the size of the loss. For these three reasons, the bank's expected losses from default are larger the longer the maturity of the forward contract.

The example also demonstrates that rolling over is an imperfect substitute to a single three-year forward contract. First, there are interim losses or gains, creating a time-value risk. For instance, the hedger does not know at what interest rates

Figure 5.3: **Spot/Forward/Money Market Diagram with bid-ask spreads**

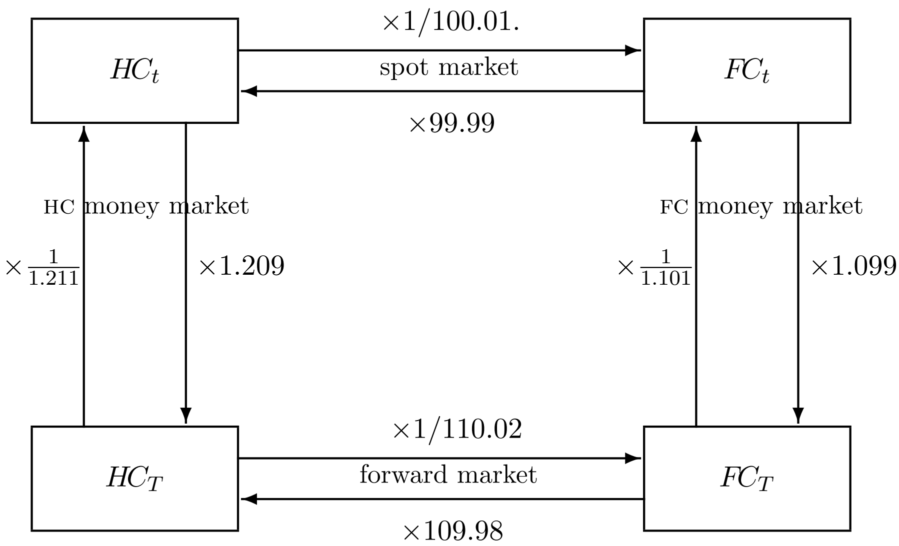

he or she will be able to finance the interim losses or invest the interim gains. Second, the hedger does not know to what extent the forward rates will deviate from the spot rates at the roll-over dates: these future forward premia depend on the (unknown) future interest rates in both currencies. Third, the total cumulative cash flow, realized by the hedger over the three consecutive contracts, depends on the time path of the spot rates between time 1 and time 3.

* * *

All this has given you enough background for a discussion of how and where forward contracts are used in practice. Among the many uses to which forward contracts may be put, the first we bring up is arbitrage, or at least the potential of arbitrage: this keeps spot, forward and interest rates in line.

5.2 Using Forward Contracts (1): Arbitrage

One question to be answered is to what extent Interest Rate Parity still holds in the presence of spreads. A useful first step in this analysis is to determine the synthetic forward rates.

5.2.1 Synthetic Forward Rates

It should not come as a surprise to you that, in the presence of spreads, the synthetic forward rates are the Worst Possible Combinations of the basic perfect-markets formula. We can immediately see this when we do the two trips on the diagram in Figure 5.3. In the diagram we started from the figures familiar from last chapter, but now setting the bids slightly below and the asks slightly above the old levels. Let's figure out the synthetic rates:

- **Synthetic bid**: the synthetic-sale trip is $FC_T \rightarrow FC_t \rightarrow HC_t \rightarrow HC_T$, and it yields

$$HC_T = FC_T \times \frac{1}{1.101} \times 99.99 \times 1.209, \tag{5.1}$$

$$\Rightarrow \text{synthetic } F_{t,T}^{bid} = \frac{HC_T}{FC_T} = 99.99\,\frac{1.209}{1.101} = 109.798. \tag{5.2}$$

- **Synthetic ask**: the synthetic-purchase trip is $HC_T \rightarrow HC_t \rightarrow FC_t \rightarrow FC_T$, and it yields

$$FC_T = HC_T \times \frac{1}{1.211} \times \frac{1}{100.01} \times 1.099, \tag{5.3}$$

$$\Rightarrow \text{synthetic } F_{t,T}^{ask} = \frac{HC_T}{FC_T} = 100.01\,\frac{1.211}{1.099} = 110.202. \tag{5.4}$$

We see that, in computing the synthetic bid rate, we retain the basic CIP formula but add the *bid* or *ask* qualifiers that generate the lowest possible combination: *bid* × *bid* / *ask*. Likewise, in computing the synthetic ask rate we pick the highest possible combination: *ask* × *ask* / *bid*. In short,

$$\text{synthetic } [F_{t,T}^{bid}, F_{t,T}^{ask}] = \left[S_t^{bid}\frac{1+r_{t,T}^{bid}}{1+r_{t,T}^{*ask}},\ S_t^{ask}\frac{1+r_{t,T}^{ask}}{1+r_{t,T}^{*bid}} \right] \tag{5.5}$$

5.2.2 Implications of Arbitrage and Shopping-around

In Figure 5.4, we illustrate the by now familiar implications of the arbitrage and shopping-around mechanisms:

1. Arbitrage ensures that the synthetic and actual quotes can never be so far apart that there is empty space between them. Thus, given the synthetic quotes 109.8-110.2, we can rule out Case 1: we would have been able to buy directly at 109.7 and sell synthetically at 109.8. Likewise, situations like Case 2 should vanish immediately (if they occur at all): we would have been able to buy synthetically at 110.2 and sell at 110.3 in the direct market.

Figure 5.4: **Synthetic and actual forward rates: some conceivable combinations**

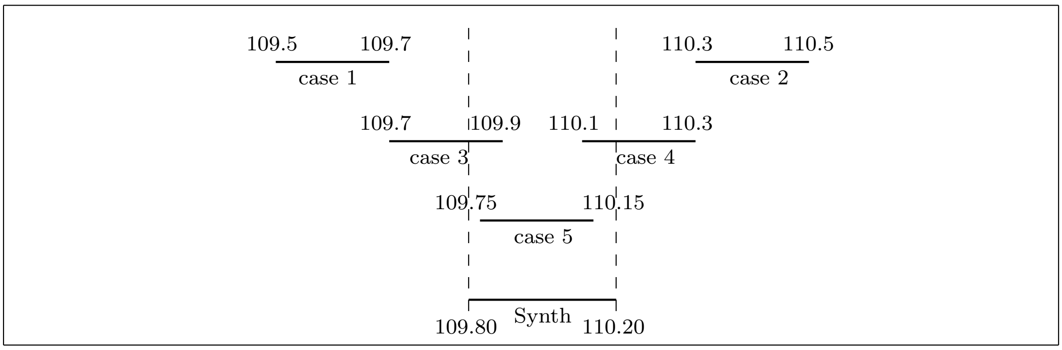

2. The usual shopping-around logic means that, in situations like Case 3 and Case 4, there would not be customers in the direct market at one side.
   - If there were only one market maker, competing against the synthetic market, Case 3 or Case 4 could occur if—and as long as—that market maker has excess inventory (Case 3) or a shortage (Case 4). These situations should alternate with Case 5.
   - But the more market makers there are, the less likely it is that not a *single* one of them would be interested in buying.[3] Likewise, with many market makers, situations where none of them wants to sell become very improbable. Thus, Cases 3 or 4 should be rare and short-lived, unless there are very few market makers.
3. With many market makers, then, Case 5 should be the typical situation: the direct markets dominates the synthetic one at both sides.

5.2.3 Back to the Second Law

How wide is the zone of admissible prices? The example has a spread of 0.4 percent between the two Worst Combinations, but that cannot be realistic at all possible maturities $T-t$. Let us first trace the ingredients behind the computations of the synthetic rates in [5.2] and [5.4]. The spot bid-ask spread is, in the example, 0.02 Pesos wide, which is about 0.02 percent. In the $(1+r)^*$ part of the formula, multiplying by 1.211 instead of 1.209 makes a difference of +0.17 percent (1.211/1.209=1.0017),

[3] In Case 3, for instance, 109.7 is by definition the best bid; all other market makers must have been quoting even lower if 109.7 is the best bid.

and the choice of $(1 + r^*)$ has an impact of +0.18 percent. Add all this up (the effect of compounding these percentages is tiny) and we get the 0.40 percent spread in the earlier calculations.

In the example, about 0.35 of this 0.40 percent comes from interest spreads. Bid-ask spreads in money markets fluctuate over time and vary across currencies, but they rise fast with time to maturity. For example, the *Wall Street Journal Europe*, January 25, 2005, mentions a Eurodollar spread of just 0.01 % *p.a.* for 30d and 0.04 % *p.a.* for 180d, implying effective spreads of less than one-tenth basis points for 30 days and 2 basis points for 180 days. So at the one-month end, interest spreads for both currencies add little to the spread between the Worst Combinations, but at 180 days most of that spread already comes from money markets. For currencies with smaller markets, spot spreads are higher but so are money-market spreads, so it is hard to come up with a general statement. Still, synthetic spreads do rise fast with time to maturity.

The widening of the spread between the Worst Combinations does give banks room to also widen the bid-ask spread on their actual quotes. As we already argued, there are good economic reasons why equilibrium spreads would go up with the horizon: markets are thinning, and the compound risk of default and exchange losses increases.[4] All this, then, explains the Second Law: banks have not only the room to widen the spreads with time to maturity but also an economic reason to do so.

This finishes our discussion of arbitrage and the Law of One Price. The second usage to which forward contracts are put is hedging, as discussed in the next section.

5.3 Using Forward Contracts (2): Hedging Contractual Exposure

The issue in this section is how to measure and hedge contractual exposure from a particular transaction. There is said to be contractual exposure when the firm has signed contracts that ensure a known in- or outflow of FC on a well-defined date. There are other exposures too, as discussed in Chapter 13; but contractual exposure is the most obvious type, and most easily hedged.

We describe how to measure the exposure from a single transaction, how to add up the contractual exposures from different contracts if these contracts mature on the same date and are denominated in the same currency and how the resulting

[4]Note that the risk is compound, a risk on a risk. The simple exchange risk under normal circumstances (i.e. assuming no failure) is hedged by closing out in the forward or, if necessary, synthetically. Exchange risk pops back up only if there is default and the bank unexpectedly needs to reverse its earlier hedge.

net transaction can be hedged. Of course, a firm typically has many contracts denominated in a given foreign currency and these contracts may have different maturity dates. In such a case, it is sometimes inefficient to hedge individually the transactions for each particular date. In Section 17.3, we show how one can define an aggregate measure of the firm's exposure to foreign-currency-denominated contracts that have different maturity dates, and how one can hedge this exposure with a single transaction.

5.3.1 Measuring Exposure from Transactions on a Particular Date

By exposure we usually mean a number that tells us by what multiple the HC value of an asset or cash flow changes when the exchange rate moves by ΔS, everything else being the same. We denote this multiple by $B^*_{t,T}$:

$$B^*_{t,T} = \frac{\Delta \tilde{V}_T}{\Delta \tilde{S}_T}. \tag{5.6}$$

Note that the delta's are for constant T—and remember that T is a known future date. That is, we are not relating a change in S over time to a change in V over time; rather, we compare two possible situations or scenarios for a future time T that differ as far as S is concerned. In continuous-math terms, we might have in mind a partial derivative. In SciFi terms, we're comparing two closely related parallel universes, each having its own S_T. Economists, more grandly, talk about comparative statics.

This is the general definition, and it may look rather other-wordly. To reassure you, in the case of contractual exposure $B^*_{t,T}$ is simply the FC value of the contract at maturity:

Example 5.5

Assume that your firm (located in the US) has an A/R next month of JPY 1m. Then, for a given change in the USD/JPY exchange rate, the impact on the USD value of the cash flows from this A/R is 1m times larger. For example, if the future exchange rate turns out to be from USD/JPY 0.0103 instead of the expected 0.0100, then the USD value of the A/R changes from USD 10,000 to 10,300. Thus, the exposure of the firm is

$$B^*_{t,T} = \frac{10,300 - 10,000}{0.0103 - 0.0100} = 1,000,000. \tag{5.7}$$

To the mathematically gifted, this must have been obvious all along: if the cash flow amounts to a known number of FC units C^*, then its HC value equals $V_T = C^* \times S_T$, implying that the derivative $\partial V_T / \partial S_T$ or the relative difference $\Delta V_T / \Delta S_T$ both equal C^*, the FC cash flow. A point to remember, though, is that while exposure *might* be a number described in a contract or found in an accounting system, it generally is not. We'll get back to this when we talk about option pricing and hedging, or operations exposure, or hedging with futures.

An ongoing firm is likely to have many contracts outstanding, with varying maturity dates and denominated in different foreign currencies. One can measure the exposure for each given future day by summing the outstanding contractual foreign-currency cash flows for a particular currency and date as illustrated in Example 5.6. Most items on the list are obvious except, perhaps, the long-term purchase and sales agreements for goods and services, with FC-denominated prices for the items bought or sold. By these we mean the contracts for goods or services that have not yet given rise to delivery and invoicing of goods and, therefore, are not yet in the accounting system. Don't forget these! More in general, contracts do not necessarily show up in the accounting system, notably when no goods have been delivered yet or no money-market transaction has been made yet.

The net sum of all of the contractual inflows and outflows then gives us the firm's net exposure—an amount of net foreign currency inflows or outflows for a particular date and currency, arising from contracts outstanding today.

Example 5.6

Suppose that a US firm, Whyran Cabels, Inc., has the following AUD commitments (where AUD is the foreign currency):
1. A/R: AUD 100,000 next month and AUD 2,200,000 two months from now
2. Expiring deposits: AUD 3,000,000 next month
3. A/P: AUD 2,300,000 next month, and AUD 1,000,000 two months from now
4. Loan due: AUD 2,300,000 two months from now

We can measure the exposure to the AUD at the one- and two-month maturities as shown in the table below:

item	30 days		60 days	
Commercial contracts in roman, financial in italic	in	out	in	out
a. A/R:	100,000	—	2,200,000	—
b. Commodity sales contracts:	0	—	0	—
c. *Expiring deposits*:	3,000,000	—	0	—
d. *Forward purchases*:	0	—	0	—
e. *inflows from forward loans in* FC:	0	—	0	—
f. A/P:	—	2,300,000	—	1,000,000
g. Commodity purchase contracts:	—	0	—	0
h. *Loan due*:	—	0	—	2,300,000
i. *Forward sales*:	—	0	—	0
j. *outflows for forward deposits in* FC	—	0	—	0
net flow	+800,000		–1,100,000	

Thus, the net exposure to the AUD one month from now is AUD 800,000 and two months from now is –AUD 1,100,000.

Note that from a contractual-exposure point of view, the future exchange rate would not matter if the net future cash flows were zero, that is, if future foreign currency denominated inflows and outflows exactly canceled each other out. This, of course, is what traditional hedging is about, where one designs a hedge whose cash flows exactly offset those from the contract being hedged. Thus, if one could match every contractual foreign currency inflow with a corresponding outflow of the same maturity and amount, then the net contractual exposure would be zero. However, perfect matching of commercial contracts (sales and purchases, as reflected in A/R and A/P and the long-term contracts) is difficult. For example, exporters often have foreign sales that vastly exceed their imports. An alternative method for avoiding contractual exposure would be to denominate all contracts in one's domestic currency. However, factors such as the counterparty's preferences, their market power, and their company policy may limit a firm's ability to denominate foreign sales and purchases in its own home currency or in a desirable third currency. Given that a firm faces contractual exposure, one needs to find out how this exposure can be hedged. Fortunately, one can use financial contracts to hedge the net contractual exposure. This is the topic of the next section.

5.3.2 Hedging Contractual Exposure from Transactions on a Particular Date

One-to-one Perfect Hedging

A company may very well dislike being exposed to exchange risk arising from contractual exposure. (Sound economic reasons for this are discussed in Chapter 12). If so, the firm could easily eliminate this exposure using the financial instruments analyzed thus far: forward contracts, loans and deposits, and spot deals. Perfect hedging means that one takes on a position that exactly offsets the existing exposure, and with contractual exposure this is easily done.

Example 5.7

We have seen, in Example 5.5, that holding a JPY T-bill with a time-T face value of JPY 1,000,000 creates an exposure of +JPY 1,000,000. Thus, to hedge this exposure, one can sell forward the amount JPY 1,000,000 for maturity T.

In the above, the purpose is just to hedge. If the firm also needs cash (in HC), it could then borrow against the future HC income from the hedge. Alternatively, the familiar spot-forward diagram tells us, one could short spot foreign exchange, that is, borrow the present value of JPY 1,000,000, and convert the proceeds into USD, the home currency. At maturity, one would then use the cash flows from the JPY T-bill to service the loan; as a result, there is no more uncommitted JPY cash left, so that no spot sale will be needed anymore, meaning that exposure is now zero.

Example 5.8

To hedge its net exposure as computed in Example 17.4, Whyran Cabels could hedge the one-month exposure with a 30-day forward sale of AUD 800,000, and the two-month exposure by a 60-day forward purchase of AUD 1,100,000.

Issue #1: Are Imperfect Hedges Worse?

Forward contracts, or FC loans and deposits allow you to hedge the exposure to exchange rates perfectly. There are alternatives. Futures may be cheaper, but are less flexible as far as amount and expiry date are concerned, thus introducing noise into the hedge; also, futures exist for heaviliy traded exchange rates only. Options are "imperfect" hedges in the sense that they do not entirely eliminate uncertainty about future cash flows; rather, as explained in Chapter 8, options remove the downside risk of an unfavorable change in the exchange rate, while leaving open the possibility of gains from a favorable move in the exchange rates. This may sound fabulous, until one remembers there will be a price to be paid, too, for that advantage.

Example 5.9

Whyran Cabels, Inc. could buy a thirty-day put option (an option to *sell* AUD 800,000 at a stated price), and a sixty-day call option (an option to *buy* AUD 1,100,000 at a stated price). Buying these options provides a lower bound or floor on the firm's inflows from the AUD 800,000 asset, and an upper bound or cap on its outflows from the AUD 1,100,000 liability.

If one is willing to accept imperfect hedging with downside risk, then one could also cross-hedge contractual exposure by offsetting a position in one currency with a position (in the opposite direction) in another currency that is highly correlated with the first. For example, a British firm that has an A/R of CAD 120,000 and an A/P of USD 100,000 may consider itself more or less hedged against contractual exposure given that, from a GBP perspective, movements in the USD and the CAD are highly correlated and the long positions roughly balance the short ones. Similarly, if an Indian firm exports goods to Euroland countries, and imports machinery from Switzerland and Sweden, there is substantial neutralisation across these currencies given that the movements in these currencies are highly correlated and the firm's positions have opposite signs.

Issue #2: Credit Risk

So far, we have limited our discussion to contractual exposure, and ignored credit risk. The risk of default, if non-trivial, creates the following dilemma:

- If you leave the foreign currency A/R unhedged (open) and the debtor does pay, you will be worse off if the exchange rate turns out to be unexpectedly low. This is just the familiar exchange risk.
- On the other hand, if you do hedge but the debtor defaults, you are still obliged to deliver foreign exchange to settle the forward contract. As soon as you hear about the default, you know that this forward contract, originally meant to be a hedge, has become an open (quasi-speculative) position. So

you probably want to *reverse* the hedge, that is, close out by adding a reverse forward.[5] But by that time the erstwhile hedge contract may have a negative value, in which case reversing the deal leads to a loss.

When there is default on the hedged FC, the lowest-risk option indeed is to *reverse* the original hedge position. For instance, if an A/R was hedged by a forward sale and if the exposure suddenly evaporates, you immediately buy the same amount for the same date. But there is about 50 percent chance that this would be at a loss, the new forward rate being above the old one. This risk, arising when a hedged exposure disappears, is called *reverse risk*.

Example 5.10

Suppose you had hedged a promised RUB 10m inflow at a forward rate of 0.033 EUR/RUB. Now you hear the customer is defaulting. So now you want to buy forward RUB 10m to neutralize the initial sale, but you soon discover that, by now, the forward rate for the same date has risen to 0.038. So if you reverse the position under these conditions, you're stuck with a loss of $10m \times (0.038 - 0.033) =$ EUR 50,000.

If the default risk is substantial, one can eliminate it, at a cost,[6] by obtaining bank guarantees or by buying insurance from private or government credit-insurance companies. Foreign trade credit insurance instruments that allow one to hedge against credit risk are discussed in Chapter 15.

Credit risk means that contractual forex flows are not necessarily risk-free. But this is just the top of an iceberg: in reality, the dividing line between contractual (or, rather, known) and risky is fuzzy and gradual in many other ways. We return to this when we discuss operations exposure, in Chapter 13.

Issue #3: Hedging of Pooled Cashflows—Interest Risk

We have already seen how one should aggregate the exposure from transactions that have the same maturity date and that are denominated in the same currency. Typically, however, a firm will have exposures with a great many different maturities. Computing and hedging the contractual exposure for each day separately is rather inefficient; rather, the treasurer would probably prefer to group the FC amounts into time buckets—say, months for horizons up to two years, quarters for horizons

[5] You could also close out with a combination of money- and spot-market deals, or negotiate an early settlement with your banker, but this necessarily produces essentially the same cashflows as those from closing out forward. Lastly, you could leave the position open until the end, and then buy spot currency to deliver as promised under the forward contract. The problem with this avenue is that the worst possible losses become bigger; so early termination of some form is usually preferred.

[6] Accountingwise this is a cost; but if the premium paid is worth the expected loss, the NPV of this deal would be low or zero.

between two and five years, and years thereafter. Then only one contract would be used to hedge the entire bucket.

Example 5.11

There are two obvious potential savings from grouping various exposures over time:

- If there are changes in sign of the flows in the bucket, *netting over time* saves money. Suppose that on day 135 you have an inflow of SEK 1.8m and the next day an outflow of SEK 1.0m. Rather than taking out two forward hedges for a total gross face value of SEK 2.8m, it would be more sensible to sell forward just SEK 0.8m for day 135, and keep the remaining SEK 1m inflow to settle the debt the next day. You'd save the extra half-spread on SEK 2m.
- *Scale economies in transaction costs.* Even if there are no changes in sign—for example, if the firm is a pure exporter—the total commission cost of doing one weekly deal of SEK 500,000 will be lower than the cost of doing five daily deals of about SEK 100,000.

One should be aware that, if pooling over time is carried too far, a degree of interest risk is introduced. Suppose, to keep things simple, that Whyran Cabels faces an inflow of SEK 100m at the beginning of year $t+5$, and one of SEK 50 at the end of that year. They could hedge this by selling forward SEK 150 dated July 1. Interest risk creeps in here because the SEK 100m that arrives on January 2 will earn interest for six months, while Whyran will have to borrow about SEK 50m because they sold forward the SEK 50m for a day predating the actual inflow. If the horizon is substantial and the potential amount of interest at play becomes nontrivial, the company can hedge the interest risk by forward deposits and loans. The example that follows assumes you know these instruments; if not, skip the example or return to Appendix 4.7 first.

Example 5.12

Suppose the forward interest rates 5×5.5 years are 3.50-3.55 % *p.a.*, and the forward interest rates 5.5×6 years are 3.75-3.80 % *p.a.*.[7] Then Whyran Cabels can do the following:

1. Arrange a deposit *ad* SEK 100m, 5 against 5.5 years, at the bid rate of 3.5 % *p.a.*, that is, 1.75 percent effective over six months. This will guarantee a SEK inflow of 101.75m on July 1st.
2. Arrange a loan with final value SEK 50m, 5.5 against 6 years at the ask rate of 3.8 % *p.a.*, that is, 1.9 percent effective over six months. The proceeds of the loan, on July 1st, will be 50m/1.019 = 49,067,713.44.
3. Sell forward the combined proceeds of the deposit (SEK 101.75m) and the loan (SEK 49.07m) for July 1.

[7] See the Appendix to Chapter 4 about forward interest rates.

Issue #4: Value Hedging versus Cash-Flow Hedging?

An extreme form of grouping occurs if the company hedges all its exposures by one single position. One simple strategy would be the following:

- Compute the PV, in forex, of all FC contracts. Call this PV_c^* (c for contract).
- Add a FC position in the bond or forward market with PV_h^* (h for hedge)
- The naive full hedge solution would then be to set $\text{PV}_h^* = -\text{PV}_c^*$.

Example 5.13

Suppose the spot interest rates are 3.4 % *p.a.* compound for 5 years, and and 3.45 % *p.a.* compound for 6 years. Then, assuming these are the company's only FC positions, Whyran Cabels can hedge its 5- and 6-year SEK debt as follows:

1. Compute $\text{PV}_c^* = 100m/1.034^5 + 50m/1.0345^6 = 125.4m$ SEK.
2. Arrange a loan with the same PV. If the loan is for one year and the one-year interest rate is 3 percent, the face value is $125.4 \times 1.03 = 129.2m$.

The reasoning behind this hedging rule is that, if the spot exchange rate moves, the effect on the PVs of the contractual position and the hedge position will balance out, thus leaving the firm's total PV unaffected. It is, however, important to realize that this argument assumes that the FC PV s of h and c are not changing, or at least that any changes in these PV *s are identical. However, foreign interest rates can change, and these shifts are likely to differ across the time-to-maturity spectrum. And even if the shifts were identical for all interest rates, the PV of the 5- and 6-year items would still change by far more than the 1-year position. Thus, PV hedging may again induce a big interest-rate risk. This is why the full hedge with just PV-matching was called "naive", above.

This can be solved by throwing in an interest-risk management program. But maturity mismatches can also lead to severe liquidity problems if short-term losses are realized while the offsetting gains remain unrealized, for the time being. A simpler solution would accordingly be to abandon the PV-hedging policy. If every single exposure is hedged by a hedge for the same date, then the impact of interest-rate changes is the same for PV_h^* and PV_c^*. This would still be approximately true if exposures are grouped into buckets that are not too wide, and if the hedge has a similar time to maturity.[8] This is why, in Example 5.12, we hedged the 5- and 6-year loan by a position at 5.5 years. In fact, since the 5-year flow is much larger

[8] Also, group inflows and outflows into separate buckets before you compute durations. (Durations for portfolios with positive and negative positions with similar times to maturity can lead to absurdly large numbers, because of leverage.) Then add a hedge on the side with the smaller PV, in absolute size.

than the 6-year one (100m versus 50m), the hedge horizon should perhaps be closer to 5 years than to six. For example, one could go for a duration-matched hedge, the one that protects the company against small, parallel shifts in the term structure:[9]

Example 5.14

Assuming the same data, Whyran Cabels can do the following:

1. Compute $\text{PV}^*_c = 100m/1.034^5 + 50m/1.0345^6 = 125.4m$ SEK
2. Compute the duration: $\frac{100m/1.034^5}{125.4m} \times \frac{5}{1.034} + \frac{50m/1.0345^6}{125.4m} \times \frac{6}{1.0345} = 5.15$ years (5 years, 54 days).
3. Arrange a loan with the same PV and duration. If 5- and 6-year rates move by the same (smallish) amount, then the effect of a shift in the term structure will equally affect the hedge instrument and the hedged positions.

As a final note, we add that complete value hedging, where the company takes one single position per currency to cover all the risks related to that currency regardless of their time to maturity, is mostly a textbook concept, even in financial companies. What does happen is hedging of net exposures that expire at dates that are close to each other; few CFOs are venturing to go any further. The complexity of the interest hedge, and the need to continuously update the interest and currency positions are obvious issues. Also, bear in mind that even if the PVs of the combined exposures and of the hedge could be kept in perfect agreement, there still is the problem that the expiry dates do not match. Cash losses may be matched by capital gains, but the latter are unrealized and unrealizable, implying that there could be serious liquidity problems. Another issue with company-value hedging is that even "contractual" exposures are never *quite* certain, as we already noted; moreover, most cashflows foreseen for a few months out are not contractual anyway, and uncertainties about non-contractual foreseen flows are often deemed to be too high to make hedging safe or reliable, to managers. We return to the issues associated with non-contractual cash flows in Chapter 13. Value hedging, in short, mainly exists in academic papers, where the managers and bankers have already read the article and therefore are as well informed as the author of the article assumes them to be. In reality, value hedging is confined to a few very simple, well-understood structures like risk-free forex positions or derivatives rather than being applied to the company as a whole.

This finishes our discussion of the second way companies and individuals use forward contracts, hedging. Later on in this book we will discuss other applications of hedging, including hedging of operating exposure (Chapter 13) and hedging for the purpose of managing and pricing of derivatives (Chapters 8, 9, and 14). The third possible application of forward contracts is speculation, as discussed in the next section.

[9]If duration is not a familiar concept, close your eyes and think of England; then skip the Example.

5.4 Using Forward Contracts (3): Speculation

What is speculation? One possible definition is that a speculator takes a position in currencies (or commodities or whatever) for purely financial reasons, not because she needs the asset or wants to hedge another position. In that sense, speculators are the agents that pick up the positive or negative net position, long minus short, left by all hedgers taken together. The forward contracts must be priced such that total net demand by hedgers and speculators is zero.

On reflection, however, almost all investments are for purely financial reasons, so by that definition almost all investors would be speculators. So while this is a perfectly valid definition, it does not necessarily match what the average person has in mind when hearing the word speculation. Many people would say that speculation involves risk-seeking, in contrast to hedging where risk is reduced rather than sought. Again, we should refine this: even buying the market portfolio involves taking risk, so by that standard most investors are again speculators. Perhaps, then, the crucial element that distinguishes speculation from ordinary investment is giving up of diversification, that is, taking positions that deviate substantially from weights chosen by the average investor in a comparable position.

If this is what we mean by speculation, the question arises whether such speculation can be rational, for risk-averse investors; shouldn't normal investors diversify rather than putting an unusual amount of money into a few assets? The answer is that speculation, or underdiversification, can be rational provided that there is a sufficient expected return that justifies giving up diversification. Extra expected returns arise from buying underpriced assets or shortselling overpriced assets. But the underdiversified speculator must realize that, by deeming some assets to be under- or overpriced, her or his opinion is necessarily in disagreement with the market's. Indeed, if the entire market had concurred that asset X is underpriced and asset Y overvalued, then you would not find any counterparts to trade at these rates, and prices would already be moving so as to eliminate the mispricing. In short, an underdiversified speculator thinks that (a) she or he spots mispricing which the market, foolishly, does not yet notice; (b) the market will soon see the error of its ways and come around to the speculator's view; and (c) the gains from that hoped-for price adjustment justify the underdiversification resulting from big positions in the mispriced assets.

In this section we discuss speculation on the spot rate, the forward rate, and the swap rate. In the examples we use speculation in the meaning of intentional underdiversification.

5.4.1 Speculating on the Future Spot Rate

Example 5.15

Suppose Milton Freedman is more optimistic about the Euro than the market. As

Figure 5.5: **Speculating in the spot market**

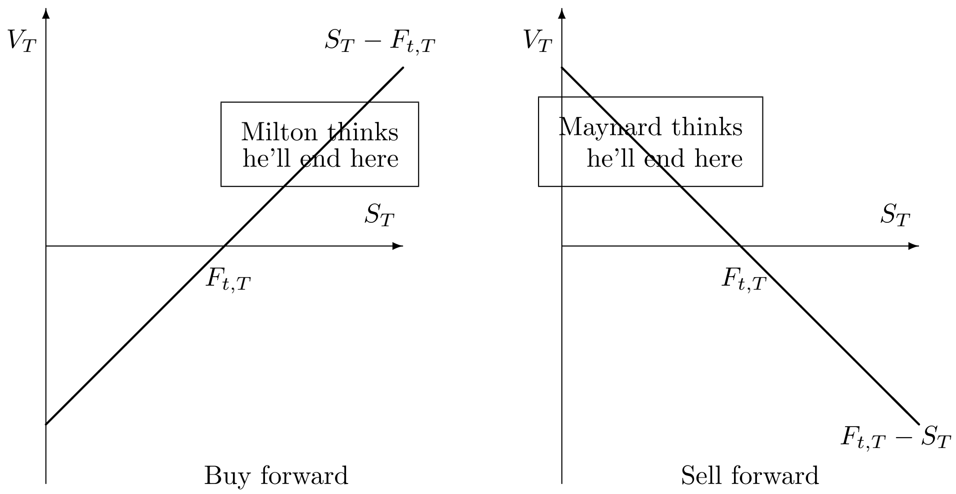

we know, the profit from buying forward will be $\tilde{S}_T - F_{t,T}$. Almost tautologically, the market thinks that the expected profit, after a bit of risk-adjustment is zero—otherwise the forward price would have moved already. But Milton thinks that, in reality, there is more of the probability mass to the right of $F_{t,T}$, and less to the left, than the market realizes. Since the potential of profits is underestimated and the room for losses overrated, Milton thinks, a forward purchase is a good deal, warranting a big position.

Example 5.16

Suppose Maynard Keenes is less optimistic about the Dollar than is the market. The profit from selling forward will be $F_{t,T} - \tilde{S}_T$ with a risk-adjusted expectation of zero—according to the market. But Maynard knows more than the market (or at least he thinks he does): depreciations are more probable, and appreciations less likely, than the market perceives. Betting on depreciations, Maynard sells forward.

In both cases, the speculator thinks that the chances of ending in the red are overrated, and the chances of making a profit underrated.[10] Note also that the forward position is closed out at the end by a spot transaction: at time T, Milton has to sell spot to realize the gain he hopefully made; and Maynard must buy spot at T because under the initial forward contract he has promised to deliver. In hedge applications, in contrast, no spot deal is needed because there already is a commercial contract which generates an in- or outflow at T.

[10] To the purists: yes, the argument is sloppy: I should talk not about chances of profits, but partial expectations. But you all know what I mean.

Of course, speculation can also be done in the spot market. Relative to buying spot, a forward purchase has the additional feature of automatic leverage: it is like buying a FC deposit already financed by a HC loan. Likewise, one alternative to selling forward is to borrow FC and sell the proceeds spot; but the extra feature in the forward sale is that the foreign currency is automatically borrowed. Here, the leverage is in FC. In either case, the leverage is good, at the private level, in the sense that positions can be bigger; but of course the risk increases correspondingly. The leverage also allows more people to speculate. This is, socially, a good thing if these extra players really do know more than the market does: then speculators are pushing prices in the right direction. And even if their opinions are, on average, no better then the other players, speculators would still help: the larger the number of people are allowed to vote on a price, the smaller the average error.

5.4.2 Speculating on the Forward Rate or on the Swap Rate

Suppose that—at time t, as usual—you want to speculate not on a future spot rate $\tilde{S}_T$ but on a future forward rate: you think that, by time T_1, the forward rate for delivery at T_2 will have gone up relative to the current level. So we speculate on $\tilde{F}_{T_1,T_2}$ instead of $\tilde{S}_{T_1}$. For example (see Figure 5.6), current time may be January and the current rate for delivery on June 1 (=T_2) may be 100.7, but you feel pretty confident that, by April 1 (=T_1), the rate for delivery early June will be higher than that. You would

- buy forward now (at t) for delivery on June 1, and
- early April, close out—that is, sell forward for June 1—at a rate that right now (in January) is still unknown.

This way, in April you will lock-in a cash flow of $\tilde{F}_{T_1,T_2} - F_{t,T_2}$, which will then be realized end June. For example, if in April the June rate turns out to be 101.6, up from 100.7, you make 101.6–100.7 = 0.9 per currency unit; or if, against your expectations, the rate falls to 100.1, you lose 0.6 per currency unit. The general net result, in short, will be $\tilde{F}_{T_1,T_2} - F_{t,T_2}$, locked in at T_1 and realized at T_2.

Of course, speculating on a drop in the forward rather than a rise works in reverse: you would sell forward now (at t) for delivery in June, and in April you would then close out and lock in the time-1 gain (?or loss?), $F_{t,T_2} - \tilde{F}_{T_1,T_2}$ to be realized at T_2.

Note that this boils down to speculation on the sum of the spot rate and the swap rate. Most of the uncertainty originates from the spot rate, however. So what would you do if you want to speculate just on the swap rate, not obscured by the spot exchange rate? And what exactly is the underlying bet?

The nature of the bet would be different. If you simply speculate on a rise in the spot rate you bet on a difference between the current (risk-adjusted) expectation and the subsequent realization. If you speculate on the future swap rate, in contrast,

Figure 5.6: **Speculating on a rise in the swap rate**

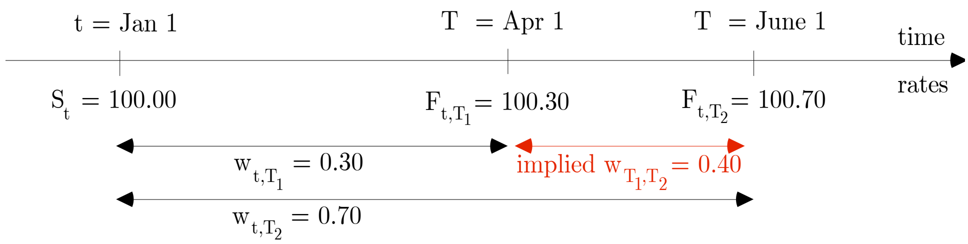

you are placing a bet on future revisions of the expectation. Consider the example in Figure 5.6. On January 1, the swap rate for delivery on April 1 is 0.30, implying that the risk-adjusted expected rise is 0.30 over that horizon. On the same date, the 6-month swap rate is 0.70, implying a risk-adjusted expected rise by 0.70 over six months. Implicit in these numbers is a risk-adjusted expected rise of 0.70–0.30=0.40 between April 1 and June 31. Suppose that you feel pretty certain that, by April 1, the market will revise its expected three-month rise upward. Your bet is that, on April 1, the three-month swap rate will exceed 0.40.

How would you do it? The answer, as we verify in the next Example, is as follows:

- you speculate on a rise of the entire forward rate (spot plus swap), as before;
- but you immediately also hedge away the spot-rate risk component, by a forward sale for delivery in April, leaving you with exposure to just the swap rate;
- you gain if, and to the extent that, the future swap rate exceeds the difference between the current swap rates (june minus April).

To explain this via an example, let us again consider a bet that the swap rate will rise:

Example 5.17

Current data:

Spot	Date T_i	Forward	swap rate
100	Apr 1	100.3	0.3
id	Jun 31	100.7	0.7
Spread Jun v Apr		0.4	0.4

The little table below lists the two ingredients in the combined strategy (the speculative bet on a fire in the forward rate, and the spot hedge) and, for each of these, the actions undertaken now and in April, plus the payoffs. The payoff of the first component is the difference between the April forward (for delivery in June) and the initial one, 100.07; the April rate is immediately written as $\tilde{S}_{\text{Apr}} + \tilde{w}_{\text{Apr-Jun}}$, where $\tilde{w}$ is the swap rate:

ingredient	action at t	action at T_1 (Apr)	payoff at expiry
bet on $F_{\text{Apr}} \uparrow$	buy forward Jun	sell forward Jun	$[\tilde{S}_{\text{Apr}} + \tilde{w}_{\text{Apr-Jun}}] - 100.7$
hedge S_{Apr}	sell forward Apr	buy spot	$100.3 - \tilde{S}_{\text{Apr}}$
Combined:	forward-forward swap "out"	spot-forward swap "in"	$\tilde{w}_{\text{Apr-Jun}} - [100.7 - 100.3]$ $= \tilde{w}_{\text{Apr-Jun}} - [0.7 - 0.3]$

We see that the ultimate profit is the realized swap rate in excess of the difference of the original ones, $0.7 - 0.3 = 0.4$.

An interesting reinterpretation is obtained if we look at the "actions" in the Example's table not row by row like we did thus far, but column by column (by date, that is).

- Start with the future actions (those planned for April). What we will do in April clearly is a spot-forward swap: we will buy spot and simultaneously sell forward. (This is called a swap "*in*" because the transaction for the nearest date, the spot one, takes us into the FC.)
- What we do right now, at t, is not unsimilar: we sell forward for one date and simultaneously buy forward for another. This is called a *forward-forward swap*, and this particular one is called "*out*" because the transaction for the nearest date is a sale, which takes us out of the FC.

Thus, instead of saying that we bet on a rise in the April forward rate and hedge the April spot component, we could equally well say that we now do a forward-forward swap, April against June, and that on April 1 we reverse this with a spot-forward swap.

5.5 Using Forward Contracts (4): Minimizing the Impact of Market Imperfections

In the previous chapter we discovered that, in perfect markets, shopping around is pointless: the two ways to achieve a given trip produce exactly the same output. Among the imperfections that we introduce in this section are (a) bid-ask spreads, (b) asymmetric taxes, (c) information asymmetries leading to inconsistent default-risk spreads, and (d) legal restrictions. Each of them make the treasurer's life far more interesting than we might had surmised in the previous chapter.

5.5.1 Shopping Around to Minimize Transaction Costs

This type of problem is easily solved using the spot/forward/money-markets diagram. A safe way to proceed is as follows:

1. Identify your current position; this is where your trip starts;

2. Identify your desired end position;
3. Calculate the outputs for each of the two routes that lead from your START to your END;
4. Choose by applying the *More Is Better* rule: more output for a given level of input is always desriable.

Example 5.18

Ms Takeshita, treasurer of the Himeji Golf & Country Club (HG&CC) often faces problems like the following:

- A foreign customer has promised a large amount of USD (= FC), but today the Club needs JPY cash to pay its workers and suppliers and does not like the exchange risk either. Should the Club borrow dollars or yens?
- The next day there are excess JPY liquidities that should be parked, risk-free. Should HG&CC go for a Yen deposit, or a swapped dollar one?
- Two days later the Club wants to earmark part of its JPY cash to settle a USD liability expiring in six months. Should they keep yens and buy forward, or move into dollars right away?
- One week later, HG&CC receives USD from a customer, and orders new irons payable in USD 180d. Should the current USD be deposited and used later on to settle the invoice?

On her laptop she sees the following data:

spot	JPY/USD 99.95 - 100.05 (spread 0.10)	180d	JPY/USD 98.88 - 99.16 (spread 0.18)
JPY, 180d	1.90 - 2.10% p.a., simple (0.95 - 1.05% effective)	USD, 180d	3.90 - 4.10% p.a., simple (1.95 - 2.05% effective)

Having taken this course, Ms Takeshita organizes the data into the familiar diagram (Figure 5.7) and sets to work. Her calculations, which take her a mere 90 seconds, are neatly summarized in Table 5.1.

Note how all computations start with one unit. The true amounts are all missing from the calculations and even from the data, thus forcing you to focus on the route. In practice, once you found the best route, you can then re-scale everything to the desired size. For instance, in application 1, if the future FC income is USD 1.235m, the output is proportionally higher too.

In this context, let me point out a frequent mistake in case of problems like application 3. Assume the liability is USD 785,235. We just found that the best

Figure 5.7: **Spot/Forward/Money Market Diagram: Ms Takeshita's data**

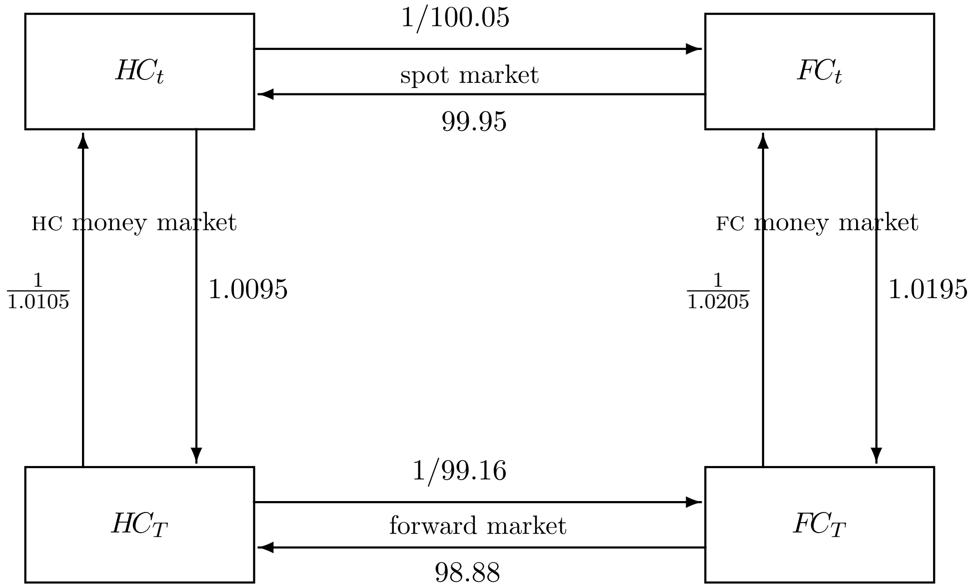

Table 5.1: **Ms Takeshita's calculations**

problem; start, end	alternatives & output
finance FC-denominated A/R (FC_T to HC_t)	* via FC_t: $\frac{1}{1.0205} \times 99.95 = 97.942185$ ♡♡ * via HC_T: $98.88 \times \frac{1}{1.0105} = 97.852548$
HC deposit (HC_t to HC_T)	* direct: 1.009500 ♡♡ * synthetic: $\frac{1}{100.05} \times 1.0195 \times 98.88 = 1.0075778$
invest in FC (HC_t to FC_T)	* via FC_t: $\frac{1}{100.05} \times 1.0195 = 0.010189905$ ♡♡ * via HC_T: $1.0095 \times \frac{1}{99.16} = 0.01018051$
park FC FC_t to FC_T	* direct: 1.0195 ♡♡ * synthetic: 99.95 x 1.0095 x $\frac{1}{99.16}$ = 1.01754260

way to move spot Yen into future Dollars is via the forward market, and the output per JPY input is 0.010189905. We easily calculate that the required investment by rescaling the whole operation, in rule-of-three style:

1. time-t input JPY 1 produces time-T output of USD 0.0010189905;
2. time-t input JPY $\frac{1}{0.0010189905}$ produces time-T output of USD 1;
3. time-t input JPY $\frac{785{,}235}{0.0010189905}$ produces time-T output of USD 785, 235.

$\Rightarrow$ (short version:)

$$\text{JPY}_t = \frac{785,235}{0.010189905} = 77,060,090. \tag{5.8}$$

This seems easy enough. What can (and often does) go wrong is that you mix up computational in- and outputs with financial in- and outputs. In computations or math, the term input refers to the data, and the term output to the result of the exercise. Financially, however, we defined input as what you feed into the financial system and output as what you get out of it. Sometimes the mathematical and the financial definitions coincide, but not always. In application 3, we exchange spot Yen for future Dollars, so the financial input is JPY_t and the output USD_T. But for the computations, the data is $\text{USD}_T = 785.235$ and the result is $\text{JPY}_t = 776,841.15$. If you're hasty, you risk thinking that the trip you need to make is from data (the mathematical input, future dollar) to result (the mathematical output, spot yen), while the actual money flow is in the other direction. Because of the mistake, you go through the graph the wrong way, using borrowing not lending rates of return and bid exchange rates instead of ask. In short, it is tempting to work back from the end point (USD_T) to the starting point: how much HC_t is needed for this? If you are really good, you will remember that going from financial output to financial input means going "against" the arrows, and chosing on the basis of a "Less Is Better" rule (less input for a given output is better). But if you are new to this, it may be safer to start by provisionally setting $HC_t = 1$, then identify the route that delivers most output (FC_T), and finally rescale the winning trip such that the end output reaches the desired level.

A second comment is that, in the second and fourth problem, the direct deposits yield more than the synthetic ones. This is what one would expect, at least if the rates are close to interbank rates. But if the problem is retail, a small FC deposit may earn substantially less than the wholesale rate (which starts at USD 1m or thereabouts), and under these circumstances the direct solution may be dominated by the indirect alternative.

Example 5.19

Suppose that the HG&CC holds a lot of JPY so that it gets interbank rates for these; but its USD deposits are small. If the rate she gets on USD were less than 3.58 % *p.a.*, Ms Takeshita would be better off moving her USD into the JPY market for six months.

On the basis of the above, one would expect that, in the wholesale market, swapping of deposits or loans should be very rare: a three-transaction trip should not be cheaper than the direct solution. But this conjecture looks at bid-ask costs only. In practice we see that swaps are often used, despite their relatively high

Table 5.2: **HC and swapped FC investments if only interest is taxed**

	Invest CLP 100		Invest NOK 1 and hedge	
initial investment		100.00	$1 \times 100 =$	100.00
final value	$100 \times 1.21 =$	121.00	$[1 \times 1.10] \times 110 =$	121.00
income		21.00		21.00
interest		21.00	$[1 \times 0.10] \times 110 =$	11.00
capgain		0	$110 - 100 =$	10.00
	Neutral taxes, 33.33%			
taxable		21.00		21.00
tax (33.33 %)		7.00		7.00
after-tax income		14.00		14.00
	Only interest is taxed, 33.33%			
taxable		21.00		11.00
tax (33.33 %)		7.00		3.67
after-tax income		14.00		17.33

transaction cost, if there is another advantage: fiscal, legal, or in terms of credit-risk spreads. We start with the tax issue.

5.5.2 Swapping for Tax Reasons

In the previous chapter we saw that swapped FC deposits and loans should yield substantially the same rate before tax, and therefore also after tax if the system is neutral. But in many countries, under personal taxation, capital gains are tax exempt and capital losses not deductible while interest income is taxed. A swapped FC deposit in a strong currency then offers an extra tax advantage: part of the income is paid out as a capital gain and is, therefore, not taxed. In Table 5.2, we go back to an example from the previous chapter and add the computations for the case where capital gains/losses are not part of taxable income. The swapped NOK deposit now offers a CLP 3.33 extra because of the tax saved on the CLP 10 capital gain.

If this is the tax rule, the implications for a deposit are as follows:

1. If the FC risk-free rate is above the domestic rate, the HC deposit does best;
2. When there are many candidate foreign currencies: the lower the FC interest rate, the higher the forward premium, so the bigger the capital gain and therefore the larger the tax advantage.

DoItYourself problem 5.2

What are the rules for a loan instead of a deposit?

You should have found that if the tax rule also holds for loans, then one would like to borrow in a weak currency, one that delivers an untaxed capital gain that is paid for, in risk-adjusted expectations terms, by a huge tax-deductable interest.

Note, finally, that there could be other tax asymmetries—for instance, capital losses being treated differently from capital gains. In that case the optimal investment rules are very different. Connoisseurs will see that in that case the tax asymmetry works like a currency option—a financial instrument whose pay-off depends on the future spot rate in different ways depending on whether S_T is above or below some critical number. To analyse this we need a different way of thinking than what we did just now.

DoItYourself problem 5.3

(For this DoItYourself assignment you do need to know the basics of option pricing.) Suppose there is a tax rule that says that corporations can deduct capital losses on long-term loans from their taxable income but they need not add capital gains to taxable income. Explain why this is different from the case above. Then show that, in this case, there always is an incentive to borrow *unhedged* FC regardless of the interest rates. (Hint: re-express the effect of this tax rule in terms of the pay-off from an option.) Finally, show that, when choosing among many FC's, you would go for the highest-volatility one, holding constant the interest rates.

5.5.3 Swapping for Information-cost Reasons

Until now we ignored credit risk. In reality even AAA borrowers pay a credit-risk spread on top of the risk-free rate. If a firm compares HC and FC borrowing, it is quite conceivable that the credit-risk spread on the FC loan is incompatible with the one on the HC alternative. For instance, if both loans are offered by the same bank, the credit analyst may have been sloppy, or may simply not have read this section of the textbook on how to translate risk spreads. Or, more seriously, the FC loan offer may originate from a foreign bank which has little information, knows it has little information, and therefore asks a stiff spread just in case.[11] The rule then is that a spot-forward swap allows the company to switch the currency of borrowing while preserving the nice spread available in another currency.

Example 5.20

Don Diego Cortes can borrow CLP for 4 years at 23 percent effective, 2 percent above the risk-free rate; and he can borrow NOK at 12 percent, also 2 percent above the

[11] Banks hate uncertainty. When they are facing an unfamiliar customer, they particularly fear *adverse selection*. That is, if the bank adds too stiff a credit-risk premium the customer will refuse, leaving the bank no worse off; but if the bank asks too little, the borrower will jump at it, leaving the bank with a bad deal. In short, unfamiliar customers too often mean bad deals.

risk-free rate. Being an avid reader of this textbook, he knows that the difference between the two riskfree rates reflect the market's opinion on the two currencies; no value is created or destroyed, everything else being the same, if one switches one risk-free loan for another, both at the risk-free rates. But the risk spreads are different: one *can* pay too much, here, and Don Diego especially feels that two percent in a strong currency (NOK) is not attractive relative to 2 percent extra on the Peso.

If, for some exogenous reason, Don Diego prefers NOK over CLP, the solution is to borrow CLP and swap into NOK:

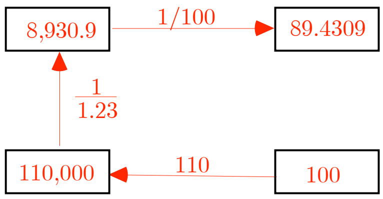

If FC_T is set at 100,000, then a direct loan at 12 percent produces $FC_t = 100,000/1.12 = 89,285.71$; but the swapped Peso loan ($FC_T \rightarrow HC_T \rightarrow HC_t \rightarrow FC_t$) yields 100,000 × 110 / 1.23 /100 = 89,430.90. Stated differently, Don Diego can borrow synthetic NOK @ (100,000 – 89,430.90)/89,430.90 = 11.81 percent instead of 12 percent.

One message is that, when comparing corporate loans in different currencies, one should look at risk spreads not total interest rates. Second, when comparing spreads we should also take into account the strength of the currency. For example, 2 percent in s strong currency is worse than 2 percent in a low one. We show, below, that the strength of the currency is adequately taken care of by comparing the PVs of the risk spreads, each computed at the currency's own risk-free rate: a 2 percent risk spread in a low-interest-rate currency then has a higher PV than a 2 percent spread in a high-rate currency. A related point, relevant for credit managers who need to translate a risk spread from HC to FC, is that two spreads are equivalent if their PVs are identical. Note that these results hold for zero-coupon loans; the version for bullet loans with annual interest follows in Chapter 7.

Example 5.21

- Don Diego can immediately note that, for the CLP alternative, the discounted spread is $0.02/1.21 = 1.65289$ percent, better then the NOK PV of $0.02/1.10 = 1.81818$ percent.
- Don Diego's banker can compute that, when quoting a NOK spread that is compatible with the 2 percent asked on CLP loans, he can ask only 1.81 percent:

$$\frac{0.02}{1.21} = \frac{0.0181}{1.10} = 0.0165289. \tag{5.9}$$

  This, as we saw before, is exactly the rate that Don Diego got when borrowing CLP and swapping.

DoItYourself problem 5.4

Here is a proof without words. Add the words—*i.e.* explain the proof to a friend who is obviously not as bright as you are. We denote the risk spreads by ρ and ρ^*, respectively.

$$\text{swapped HC loan yields } \begin{array}{c}\text{more,}\\ \text{the same, if}\\ \text{less,}\end{array} \quad \frac{F_{t,T}}{1+r+\rho} \times \frac{1}{S_t} \gtreqless \frac{1}{1+r^*+\rho^*}$$
$$\Updownarrow$$
$$\frac{1+r}{1+r^*}\,\frac{1}{1+r+\rho} \gtreqless \frac{1}{1+r^*+\rho^*}$$
$$\Updownarrow$$
$$\frac{1+r}{1+r+\rho} \gtreqless \frac{1+r^*}{1+r^*+\rho^*}$$
$$\Updownarrow$$
$$\frac{1+r+\rho}{1+r} \lesseqgtr \frac{1+r^*+\rho^*}{1+r^*}$$
$$\Updownarrow$$
$$\frac{\rho}{1+r} \lesseqgtr \frac{\rho^*}{1+r^*}. \qquad (5.10)$$

5.5.4 Swapping for Legal Reasons: Replicating Back-to-Back Loans

In the examples thus far, we used the swap to change the effective denomination of a deposit or a loan. We now discuss reasons to work with a stand-alone swap. The main use of this contract is that it offers all the features of back-to-back loans (that is, two mutual loans that serve as security for each other), but without mentioning the words loan, interest, or security. We proceed in three steps. First we explain when and why back-to-back loans may make sense. We then establish, via an example, the economic equivalence of a swap and two back-to-back loans. Lastly we list the legal advantages from choosing the swap representation of the contract over the direct back-to-back loan.

Why Back-to-back Loans may make Sense

The most obvious reason for a back-to-back-like structure is providing security to the lender.

Example 5.22

During the Bretton Woods period (1945–1972), central banks often extended loans to each other. For example, to support the GBP exchange rate, the Bank of England (BoE) would buy GBP and sell USD. On occasion it would run out of USD. Hoping that the pressure on the GBP (and the corresponding scarcity of USD reserves) was temporary, the BoE would borrow USD from, say, the Bundesbank (Buba), the central bank of Germany. The Buba would ask for some form of security for such a loan. In a classical short-term swap deal, the guarantee was in the form of an equivalent amount of GBP to be deposited with the Buba by the BoE. Barring default, on the expiration day the USD and the GBP would each be returned, with interest, to the respective owners. If either party would default, the other was automatically exonerated of its own obligations and could sue the defaulting party for any remaining losses.

Example 5.23

The central bank of the former Soviet Union often used gold as security for hard-currency loans obtained from western banks, but repeatedly failed to pay back the loans. For the western counterparty, the risk was limited to the face value of the loan minus the market value of the gold. The Soviet Union always made good this loss.

Example 5.24

Companies often post bonds or T-bills or other tradable securities as guarantee to a loan. One way to view this is that the borrower lends the bonds to the bank, which in return then lends money to the company. The bank can confiscate the bonds and sell them off if the company fails to pay back the loan.

Other applications are of the pure back-to-back loan type: a customer lends money to the bank, which in turn lends back money to the customer and uses the deposit as security for the loan. One motivation may be money laundering:

Example 5.25

After a long and successful career in the speak-easy business, Al-C wants to retire and spend his hard-won wealth at leisure. Fearing questions from the tax authorities, he deposits his money in the Jamaica office of a big bank, and then borrows back the same amount from the NY office of that bank. The deposit serves as security for the loan: if Al is unexpectedly *taken out*, the bank confiscates the deposit in lieu of repayment of the loan. And when questioned by the tax inspectors as to the source of the money he spends so freely, Al can prove it is all borrowed money.[12]

Another motivation is avoidance of exchange restrictions or other costs of moving money across borders. Back-to-back loans (or parallell loans) were often inspired by the investment dollar premium that existed in the UK from the late sixties to

[12] The example lacks credibility because the taxman's next question would be *why* the bank lent so much to Al. So this can only be done on a small scale, by persons or companies that could have borrowed such amounts without the guaranteeing deposit.

Figure 5.8: **The Parallell loan – Example 1**

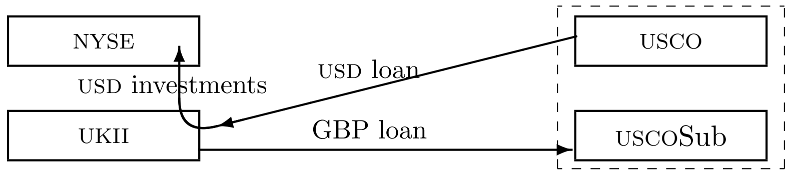

the mid-seventies and made it expensive for companies to buy dollars for foreign investments.[13] Back-to-back loans, promoted and arranged by UK merchant banks, were a way to avoid this investment dollar premium.

Example 5.26

Suppose a UK institutional investor (UKII) would like to invest in the NYSE. The trick is to find a foreign company (say, USCO) that wants to extend a loan to its UK subsidiary. The USCO, rather than lending to its UK subsidiary, lends USD to UKII. Thus, the UKII borrows USD and pays them back later, which means that it does not have to buy USD initially and that there is no subsequent sale of USD. In short, the investment dollar premium is avoided. The second leg of the contract is that UKII lends GBP to USCO's subsidiary, so that USCO's objectives are also satisfied. The expected gains from avoiding the implicit tax can then be divided among the parties. The flow of the principal amounts of the reciprocal loans is shown in Figure 5.8.

As it stands, the design of the back-to-back loan would be perfect if there were no default risk. Suppose, however, that USCO's subsidiary defaults on its GBP loan from the UKII. If no precautions had been taken, UKII would still have to service the USD loan from USCO, even though USCO's subsidiary did not pay back its own loans. Writing a right-of-offset clause into each of the separate loan contracts solves this problem. If USCO's subsidiary defaults, then UKII can suspend its payments to USCO, and sue for its remaining losses (if any)—and vice versa, of course. Thus, the right of offset in the back-to-back loan is one element that makes this contract similar to mutually secured loans. The similarity becomes even stronger if you consolidate

[13] In those years, the UK had a two-tier exchange rate. Commercial USD's (for payments on current account, like international trade and insurance fees) were available without constraints, but financial USD's (for investment) were rationed and auctioned off at premiums above the commercial rate. These premiums, of course, varied over time and thus were an additional source of risk to investors. In addition, the law said that when repatriating USD investments, a UK investor had to sell 25 percent of his financial USD in the commercial market; the premium lost was an additional tax on foreign investment. In summary, there was quite a cost attached to foreign investment by UK investors.

Figure 5.9: **The Parallell loan – Example 2**

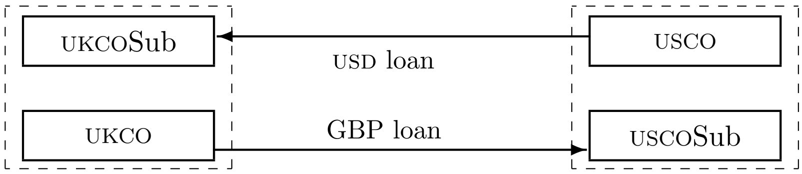

USCO with its subsidiary and view them as economically a single entity—see the dashed-line box in the figure. Then, there clearly is a reciprocal loan between USCO and UKII, with a right of offset.

Example 5.27

If USCO also faced capital export controls (for example, Nixon's "voluntary" and, later, mandatory controls on foreign direct investment), there would be no way to export USD to the UK counterpart. Suppose that there also was a UK multinational that wanted to lend money to its US subsidiary, if it were not for the cost of the investment dollar premium. The parallel loan solves these companies' joint problem, as shown in Figure 5.9. (The diagram shows the direction of the initial principal amounts.) USCO lends UKCO dollars in the US, without exporting a dime, while UKCO lends pounds to USCO's subsidiary in the UK (and, therefore, is making no foreign investment either).

Thus, no money crosses borders, but each firm has achieved its goal. UKCO's subsidiary has obtained USD, and USCO's subsidiary has obtained GBP, and the parents have financed the capital injections. This parallel loan replicates the reciprocal loan inherent in the short-term swap when we consolidate the parents with their subsidiaries (see the dashed boxes). In addition, the parallel loan typically has a right-of-offset clause that limits the potential losses if one of the parties defaults on its obligations.

Example 5.28

Suppose you have left Zimbabwe, where you lived most of your life, but you are not allowed to take out the Zimbabwe Dollars you accumulated during your career. What you can do is try to find someone who, puzzlingly, wants to invest money in Zimbabwe, and to convince that party to lend his pounds to you in London, while you undertake to finance his Zimbabwe investment. (One occasionally sees such proposals in the small-ad sections of *The Times* or *The Economist*.) Both parties would feel far safer if there also is a right-of-offset clause in the loans.

Now that we understand why people may want mutually secured loans, we turn to the link between these contracts and swaps.

The Economic Equivalence between Back-to-back Loans and Spot-forward Swaps

Let us go back to the USD loan from Buba to BoE, and let's add some specific figures. We then summarise the contract in a table:

Example 5.29

The little table below shows this deal from BoE'spoint of view: the USD loan in the second column, the GBP deposit in the third. (Ignore the fourth column for the time being.) The rows show for each contract the promised payments at t and T, assuming a dollar loan of 100m, a spot rate of USD/GBP 2.5, and an effective six-month rate of 3% on dollars and 5% on pounds. Outflows, from BoE's point of view, are indicated by the $< >$ signs around the amounts.

	USD 100m borro'd at 3 percent	GBP 40m lent at 5 percent	
t	USD 100.0m	<GBP 40.0m>	= spot purchase of USD 100m at 2.5
T	<USD 103.0m>	GBP 42.0m	= forward sale of USD 103m @ 2.4523

The funny thing is that if one looks at the table by date (*i.e.* row by row) rather than by contract (*i.e.* column by column), one sees for date t a spot conversion of USD 100m into GBP, at the spot rate of 2.5. For the end date, there is a promised exchange of USD 103m for GBP 42m, which sounds very much like a forward deal. Even the implied forward rate is the normal forward rate, as one can see by tracing back the numbers behind that rate:

$$2.4523 = \frac{103}{42} = \frac{100}{40}\frac{1.03}{1.05} = S_t \frac{1 + r_{\text{USD}}}{1 + r_{\text{GBP}}} = F \tag{5.11}$$

Thus, depending on one's preferences, the promised cash flows can be laid down either in two loan contracts that serve as security for each other, or in a spot contract plus an inverse forward deal—a spot-forward swap. But the similarity goes beyond the promised cashflows: even in the event of default the two stories still have the same implication. If, say, BoE defaults, then under the two-loans legal structure Buba will invoke the security clause, sell the promised GBP 42m in the market rather than give them to BoE, and sue if there is any remaining loss. Under the swap contract, if BoE defaults, Buba will invoke the right-of-offset clause, sell the promised GBP 42m in the market rather than give them to BoE, and sue if there is any remaining loss. Thus, the two contract structures are, economically, perfect substitutes. But lawyers see lots of legal differences, and many of these make the swap version more attractive than the mutual-loan version.

Legal Advantages of the Swap Contract

Simplicity. Legally speaking, structuring the contract as a spot-forward transaction is simpler than the double loan contract described earlier.

Figure 5.10: **A Bank betraying its Pawnshop Roots**

Source Author's picture

Example 5.30
In a repurchase order (repo) or repurchase agreement, an investor in need of short-term financing sells low-risk assets (like T-bills) to a lender, and buys them back under a short-term forward contract. This is another example of a swap contract (a spot sale reversed in the forward market). In terms of cash flows, this is equivalent to taking out a secured loan. Because of the virtual absence of risk, the interest rate implicit between the spot and forward price is lower than an ordinary offer rate and differs from the lending rate by a very small spread, called the bank's "haircut." In the case of default, the bank's situation is quite comfortable because it is already legally the owner of the T-bills.

Repo lending is a fancy name for what is done in pawnshops. In fact, banking and pawning used to be one and the same. In Germany, a repo is called a *Lombard* (and the repo rate is called the Lombard rate), after the North-Italian bankers that introduced such lending in the Renaissance times; in Dutch, *lommerd* just means Pawnshop. The Catholic Church, incensed at the high rates charged, then started its own Lombard houses with more reasonable rates. These institutions were often called *Mons Pietatis*, Mount(ain) of Mercy; some still exist nowadays and a few have grown into big modern banks. The oldest surviving bank, Monte dei Paschi de Siena (1472), is one of these. Figure 5.10 shows a Spanish example.

We know, from Chapter 2, that central banks can steer the money supply upward by lending money to commercial banks, or downward by refusing to roll over old loans to banks. Nowadays these loans typically take the form of repo's. In many countries the repo rate has become the main beacon for short-term interest rates.

In short, simplicity and efficiency is one advantage of a swap contract over a secured or back-to-back loan. To lawyers, who do not necessarily view simplicity as a plus, the main attractions are that the words security, interest, and loan/deposit are not mentioned at all.

The term *security* is not used. If the contract involves private firms rather than two central banks, the firm's shareholders need not be explicitly informed about the implicit right-of-offset clause in a swap because a forward contract is not even in the balance sheet (see below). In contrast, if there had been two loans, the financial statements would have had to contain explicit warnings about the mutual-security clause.[14] In some countries, the clause must even be officially registered with the commercial court or so. Providing security may also be contractually forbidden if the company has already issued bonds or taken up loans with the status of *senior* bonds or loans: giving new security would then weaken the position of the existing senior claimants. Bond covenants may also restrict the firm's ability to provide new security. All these problems are avoided by choosing the swap version of the contract.

The term *Interest* is not used. Similarly, also the word interest is never mentioned in a swap contract; there is only an implied capital gain. This can be useful for tax purposes, as we saw before. In the example below, the reason is religious objections against interest.

Example 5.31

In the Middle Ages, the Catholic Church prohibited the payment of interest; swap-like contracts were used to disguise loans. Eldridge and Maltby (1991) describe a three-year forward sale for wool, signed in 1276 between the Cistercian Abbey of St. Mary of the Fountains (N-England) and a Florentine merchant. The big "margin" deposited by the merchant was, in fact, a disguised loan to the Abbey, serviced by the deliveries of wool later on. The forward prices were not stated explicitly, because the implied interest would have been made too easy to spot.

The term *loan* or *deposit* is not used. A parallel loan would have shown up on both the asset and liability sides of the balance sheet. In contrast, a forward deal is off-balance-sheet.[15] This has several advantages: (*i*) it does not inflate debt, so it leaves unaffected the debt/equity ratio or other measures of leverage; (*ii*) it does not inflate total assets, so it leaves unaffected the profit/total-assets ratio. Under the old BIS rules ("Basel I"), capital requirements on swaps were less exacting than those on separate loans and deposits (see box in Figure 5.11).

[14] Anybody involved with the firm has the right to know what assets have been pledged as security: this would mean that the firm's asset are of no use to the ordinary claimant if and when the firm defaults on its obligations.

[15] This accounting rule is not unreasonable. There is indeed a difference between a swap and two separate contracts (one asset and one liability). In the case of the swap, default on the liability wipes out the asset. For that reason, accountants think it would be misleading to show the swap contract as if it consisted of a standard separate asset and liability. The inconsistency is, however, that once an asset has been pledged as security, it remains on the balance sheet *except* for forwards, futures, swaps etc.

Figure 5.11: **Capital Adequacy Rules v 1.0 (Basel 1)**

The Bank for International Settlements of Basel, Switzerland, has no power to impose rules on banks anywhere. However, the BIS deserves credit for bringing together the regulatory bodies from most OECD countries in a committee called the BIS Committee, or the Basel Committee, or the Cooke Committee (after the committee's chairman), to create a common set of rules. The objective of establishing a common set of rules was to level the field for fair competition.

Under the original agreement the general capital requirement was 8 percent, meaning that the bank's long-term funding had to be at least 8 percent of its assets. For some assets and for off-balance-sheet positions with a right of offset, the risk was deemed to be less than the risk of a standard loan to a company, and the capital ratio was lowered correspondingly. For instance, a loan to any (!?) government or bank was assumed to have zero credit risk, and did not require any long-term capital. The rule was crude but was deemed to be better than no rule at all.

This is now called Basel 1. The more recent Basel-2 rules have replaced the 8% rule for credit risks by a system of ratings—external whenever possible, internal otherwise—and have added Value at Risk (Chapter 13) to cover market risks.

A more shady application of disguising one's lending and borrowing arose when a finance minister decided to speculate with the taxpayers' money, and used swaps for the purpose:

Example 5.32

At one EC Council meeting in the mid-1980s, even Margaret Thatcher, caught off guard, was provoked into saying that she could not 100% exclude that the UK might never ever think of discussing the option of joining a common European currency. Belgium's then finance minister, Mr Maystadt, concluded that the advent of the common currency was a matter of a few years and that it would be introduced at the official parities, without any interim realignments. From these views—which, it later turned out, were both wrong—it followed that the huge interest differential between Lira and Marks had become virtually an arbitrage opportunity. Thus, speculation was justified: one should borrow in a low-interest currency, like DEM, and invest the proceeds in a high-interest one, like ITL (the "carry trade"). Still, the country's rule-books stated that the Finance Ministry could borrow only to finance the state's budget deficit. The minister therefore signed a huge long-term swap contract instead, arguing that since the law did not mention swaps, their use was unrestricted.

The whole deal blew up in his face when the ERM collapsed in 1992 and the ITL lost one third of its value.

This has brought us to the end of our list of possible uses of forward contracts. We close the chapter with a related management application, where we are not strictly using the forward contract but rather the forward rate as a useful piece of information, notably in the case of valuation for management accounting purposes. This is discussed in the next section.[16]

5.6 Using the Forward Rate in Commercial, Financial and Accounting Decisions

5.6.1 The Forward Rate as the Intelligent Accountant's Guide

Suppose a Canadian exporter sells goods in New Zealand, on a NZD 2,5m invoice. This transaction has to be entered into the accounts,[17] and as the exporter's books are CAD-based, the accountants need to translate the amount into CAD. In this context, many accountants fall for the following fallacy: "if we sell for NZD 2.5m worth of goods, and one NZD is worth CAD 0.9, then we sell CAD 2.25m worth of goods." So these accountants would naturally use the spot rate to convert FC A/R or A/P into HC.

Why is this called a fallacy? What's wrong with the argument is that it is glossing over timing issues. True, if today we sell our wares *and* get paid second working day *and* we already convert the NZD spot into CAD right now, we'll get CAD 2.25m in our bank account on day $t + 2$. But almost all real-world deals involve a credit period. So the above story should be modified: today we sell, and we will receive NZD 2.5m in, say, 45 days. At what rate we will convert this amount into CAD depends on whether we sell forward or not. This is how a finance person worth her salt would think:

- If we do sell forward, then it would look natural to book the invoice at the forward-based value. After all, if we sell for NZD 2.5m worth of goods, and we know we'll receive CAD 0.88 per NZD, one would logically book this at CAD $2.5 \times 0.88 = 2.2$m.

- If we do not sell forward, we do not know yet what the exact CAD proceeds will be. So we have to settle for some kind of expected value or equivalent value, for the time being. Since we know that hedging does not change the economic value (at the moment of hedging, at least), we should use the same valuation procedure as if we had hedged—the forward rate, that is. So we still book this as a CAD 2.20m sale even if there is no hedging.

[16] Of course there are more exchange-rate related issues in accounting than what we discuss here, but they are not directly related to the forward rate; we relegate those to Chapter 13.

[17] In traditional accounting this is done as soon as the invoice has been sent or received. Under IFRS, this can be done as soon as there is a firm commitment. More precisely, the firm commitment is then entered at initially a zero value but can and must be updated when the invoice arrives or leaves and at any intervening reporting date. See Chapter 13 for more.

Many accountants would howl in protest. For instance, they might say, if one converts the NZD 2.5m at the forward rate, then the CAD accounting entry would depend on whether the credit period is 30 days or 60 or 90 etc. This is true. But there is nothing very wrong with it. The root of this problem is that accountants are *always* booking face values, not corrected in any way for time value. If they had used PVs everywhere, nobody would have a problem with the finding that an invoice's present value depends on how long one has to wait for the money.

This, of course, might be hard to grasp for some of the accountants. If so, at this point you take advantage of his confusion and ask him whether, if valuation for reporting purposes is done at the spot rate, there is a way to actually lock in that accounting value—that is, make sure you actually get the book value of CAD 2.25m. The only truthful answer of course is that there is no way to do this. You can then subtly point out that there *is* a way to lock in the accounting value of 2.20m: hedge forward. Giving no quarter, you then ask whether the spot rate takes into account expected exchange-rate changes and risks. Of course not, the accountant would bristle: in Accounting, there surely is no room for subjective items like "expectations" and "risk adjustments". The spot rate, he would add, is objective, as any valuation standard should be. You can then subtly point out that the forward rate actually is the risk-adjusted expectation, and that it is a market-set number not a subjective opinion. At this point your scorecard for the competing translation procedures looks as follows:

criterion	convert at S_t	convert at $F_{t,T}$
can be locked in at no cost?	no	yes
takes into account expected changes?	no	yes
takes into account risks?	no	yes
objective?	yes	yes
understandable to accountants?	yes	hm

The accountant's last stand might be that valuation at 0.88 instead of 0.90 lowers sales and therefore profits; and more profits is good. This is an easy one. First, for other currencies there might be a forward premium rather than a discount; and for A/P a discount would increase operating income rather than decreasing it. So there is no general rule as to which valuation approach would favor sales and lower costs. Second, you point out, total profits are unaffected by the valuation rule: the only thing that is affected is the way profits are split up into operating income and financial items.

Example 5.33

Suppose, for instance, that our Canadian firm does not hedge the NZD 2.5m, and at T the spot rate turns out to be 0.92. Suppose also that the cost of goods sold is CAD 1.5m. Then profits amount to 2.5m $\times$ 0.92 – 1.5m = 2.3m – 1.5m = 0.8m regardless of what you did with the A/R.

True, the operating profit does depend on the initial valuation of the A/R, but there is an offsetting effect in the capital gain/loss when the accounting value is confronted with the amount actually received:[18]

	using $S_t = 0.90$		using $F_t = 0.88$	
• at t:				
A/R	2,250		2,200	
COGS		1,500		1,500
operating income		750		700
• at T:				
bank	2,300		2,300	
A/R		2,250		2,200
capital gain/loss		50		100

5.6.2 The Forward Rate as the Intelligent Salesperson's Guide

For similar reasons, the forward rate should also be used as the planning equivalent in commercial decisions. Let us use the same data as before, except that the production cost is 2,210. If the "spot" valuation convention is followed, a neophyte sales officer may think that this is a profitable deal. It isn't: the Equivalent HC amount of NZD 2.5m is $2,5 \times 0.88 = 2,200$, not 2,250 as the spot translation would seem to have implied.

Some cerebrally under-endowed employees may think that the valuation difference is the cost of hedging, but you should know better by now. The acid test again is that the value 2,200 can be locked in at no cost, while you would have had to pay about 50 (minus a small PV-ing correction) for a non-standard forward contract (sell NZD 2,500 at 0.90 instead of at the market forward rate, 0.88). That is, locking in a value of 2,250 would cost you 50 at T, implying that the true future value is 2,200.

5.6.3 The Forward Rate as the Intelligent CFO's Guide

Lastly, in taking financing decisions we can always use the forward rate to produce certainty equivalents for FC-denominated service payments. The principle has been explained before. The CEQ idea or, equivalently, the zero-initial-value property of a forward deal, imply that no value is added or lost by replacing a loan by another one in a different currency.

[18]Note that while what I show below looks like accounting entries to the untrained eye, it violates all kinds of accounting rules and conventions. For instance, one does not immediately calculate and recognize the profit when a sale is made. Still, you can interpret it as a CEO's secret private calculations of profits and losses from this transaction; and it does convey the gist of what accountants ultimately do with this deal.

Two remarks are in order. First, the above statement ignores credit risks, as we have shown: while no value is gained/lost when adding a swap, value *is* gained when an unnecessarily high risk spread is replaced by a better one. We also should look at various fees and transaction costs, and possible non-neutralities in the tax law. All this issues make the CFO's life far more interesting than it would have been in a perfect world. Second, when stressing the CEQ property, we also assume that the market knows what it is doing. Some CFOs may disagree, or at least disagree some of the time, and turn to speculation. Others may agree that the market rates are fair but can still have a preference for a FC loan, for instance because it hedges other FC income. So even if in terms of market values nothing would be gained or lost there can still be a preference for a particular currency.

But when swaps are possible, the ultimate currency of borrowing can be separated from the currency in which the original bank loan is taken up. Thus, we first choose on the basis of costs. Then we ask the question whether the currency of the cheapest loan is also the currency we desire to borrow in. If so, then we're happy already. If not, then (*i*) a cheap HC loan can be swapped into FC if desired—e.g. to hedge other income or to speculate; or (*ii*) a cheap FC loan can be hedged, if desired. Thus, in the presence of swap and forward markets it is always useful to split the discussion of, say, what currency to borrow from what bank, into two parts: (*i*) what are the various transaction costs, risk spreads, and tax effects?; and (*ii*) do we want to change the currency of lowest-cost solution by adding a swap or a forward?

How would we sum up costs and spreads and so on? Here's an example. In the calculations we calculate all costs in PV terms, using the risk-free rate of the appropriate currency.[19]

Example 5.34

Suppose you have three offers for a loan, one year. You need EUR 1m or, at S_t = 1.333, USD 1.333m if you borrow USD. Below, I list the asked interest rate, stated as swap plus spread, and the upfront fee on the loan—a fixed amount and a percentage cost. How would you chose?
– Bank A: EUR at 3% (LIBOR) + 1.0%; upfront EUR 1000+0.50%
– Bank B: EUR at 3% (LIBOR) + 0.5%; upfront EUR 2000+0.75%
– Bank C: USD at 4% (LIBOR) + 0.9%; upfront USD 1000+0.50%

The computations are straightforward:

[19]Discounting at the risk-free rate is not 100% correct: when we want to find the PV, to the borrower or lender, of a series of payments we should take a rate that includes default risk. (The procedure with discounting at the risk-free rate, above, was derived to find equivalent payment streams from the swap dealer's point of view, who has a much safer position than the lender.) But in the presence of upfront fees it is no longer very obvious what the rate on the loan is, and the error from using the swap rate instead is small. A more in-depth discussion follows in Chapter 16.

	amount	PV risk spread	upfront	total
A	EUR 1m	$\frac{1m\times 0.010}{1.03} = 9{,}708.7$	$1000 + 5{,}000.0 = 6{,}000$	15,708.7
B	EUR 1m	$\frac{1m\times 0.005}{1.03} = 4{,}854.4$	$2000 + 7{,}500.0 = 9{,}500$	14,354.4
C	USD 1.333m	$\frac{1.333m\times 0.009}{1.04}/1.333 = 8.653.8$	$\frac{1{,}000+1.333m\times 0.005}{1.333} = 5{,}750.2$	14,404.0

So the second loan is best. The issue of whether or not to speculate then boils down to whether you are keen on selling a large amount of USD 360 days—for instance to speculate on a falling USD, or to hedge other USD income.

5.7 CFO's Summary

This concluding section has two distinct parts. First I want to simply review the main ideas you should remember from this chapter. The second item is a very much bird's view of the currency markets and their players.

5.7.1 Key Ideas for Arbitrageurs, Hedgers, and Speculators

We opened this chapter with a discussion of bid-ask spreads. Any transaction or sequence of transactions ("trip") that is not a round-trip (not a pure arbitrage transaction) can still be made through two different routes. In imperfect markets—and, notably, with positive spreads—it is a near certainty that one route will be cheaper than the other, and therefore, it generally pays to compare the two ways of implementing a "trip." The route chosen matters because, with spreads, it is mathematically impossible that for every single trip the two routes end up with exactly the same result. Equality of outcomes may hold, by a fluke, for at most one trip. And even if the difference between the outcomes of the two routes is small in the wholesale market, that difference can be more important in the retail market, where costs are invariably higher.

But there is more to be taken into consideration than spreads. Differential taxation of capital gains/losses and interest income/cost provides another reason why two routes are likely to produce different outcomes. For most corporate transactions, however, taxes may not matter, since interest and short-term capital gains (like forward premia received or paid) typically receive the same tax treatment. Lastly, information asymmetries can induce incompatibilities between the risk spreads asked by different banks; and if the loans also differ by currency, one can go for the best spread and then switch to the most attractive currency via a swap. Recall that the attractiveness of a loan is mainly determined by its (PV'ed) risk spread, not the total interest rate.

A second implication of bid-ask spreads relates to the cost of hedging. In Chapter 4, we argued that, in perfect markets, hedging has no impact on the value of the firm unless it affects the firm's operating decisions. In the presence of spreads, this needs a minor qualification, though. If a firm keeps a net foreign exchange position open, it will have to pay transactions costs on the spot sale of these funds, when the position expires. If the firm does hedge, in contrast, it will have to pay the cost in the forward market. Since spreads in the forward markets are higher, the extra cost represents the cost of the hedging operation. But we know that the cost of a single transaction can be approximated as half the difference between the bid-ask spread, so the cost of hedging is the extra half-spread, which at short maturities remains in the order of a fraction of one percent. Not zero, in short, but surely not prohibitive.

Forward contracts are often used as a hedge. Remember that there may be an alternative hedge, especially if the hedge is combined with a loan or deposit. Also, show some restraint when a single contract is to be used for hedging many exposures pooled over a wide time horizon. An extreme strategy is to hedge all exposures, duly PV'ed, by one hedge. Such a strategy involves interest risk and may also cause severe liquidity problems if the gains are unrealized while the losses are to be settled in immediate cash. It is safer and simpler to stay reasonably close to the matching of cash flows rather than hedging the entire exposed present value via a single contract.

Speculation is a third possible application. Recall that, as an underdiversified speculator, you implicitly pretend to be cleverer than the market as a whole (which, if true, probably means that reading this book is a waste of time). Speculation can be done on the spot rate, the forward rate, or the difference of the two, the swap rate. One can execute this last strategy by forward-forward and spot-forward swaps, but upon scrutiny this turns out to be just speculation on the forward rate, with the spot-rate component in that forward rate simply hedged away.

Swaps can also offer the same advantages as secured loans or back-to-back loans with, in addition, all the legal advantages of never mentioning the words security, interest, or loan. They have been the fastest-growing section of the exchange market since their emergence from semi-obscurity in the 1980s. We return to the modern currency swap in Chapter 7.

Lastly, it is recommended that you use forward rates to value contractual obligations expressed in FC. Standard practice is to use the current spot rate, but there is no way to lock in the current spot rate for a future payment; relatedly, that spot rate is not the risk-adjusted expectation or certainty equivalent of the future spot rate either. But remember that total profits are unaffected: the only impact is on the division of profits into operational v. financial income. So as long as you remember that a premium or discount is not the cost of hedging in any economically meaningful way, little harm is done by using the wrong rate.

This ends the "review" part of this concluding section. At this stage you know enough about spot and forward markets to understand the global picture. Let us consider this, too.

5.7.2 The Economic Roles of Arbitrageurs, Hedgers, and Speculators

This is the second of two chapters on forward markets. One thing you do remember from these, it is hoped, is the fact that spot, money, and forward markets are one intertwined cluster. Traditionally, players in these markets are categorized as hedgers, speculators, or arbitrageurs. For current purposes, we shall define speculation widely, including all pure financial deals, whether they are based on perceived mispricing or not. Likewise, let's temporarily broaden arbitrage to include not just strict arbitrage but also shopping-around: both help enforcing the Law of One Price. Let's now see how these markets and these players interact to arrive at an equilibrium.

The role of hedgers is obvious. In agricultural markets, for instance, soy farmers want to have some certainty about the sales value of their next crop, so they sell forward part or all of the expected harvest. Manufacturers that need soy as inputs likewise are interested in some certainty about their costs and could buy forward. Similarly, in currency markets, companies with long positions want to sell forward, and players with short positions want to buy. But if hedgers were the only players, the market might often be pretty thin, implying that the market-clearing price could occasionally be rather weird. That is where speculators and arbitrageurs come in.

The role of arbitrageurs, notably, is to make sure that a shock in one market gets immediately spread over all related markets, thus dampening its impact. For instance, if excess sales by hedgers would require a sharp drop in the forward rate to clear the market, then CIP means that the spot rate will feel the pressure too; and if the spot rate moves, all other forward rates start adjusting too. What happens, in principle, is that arbitrageurs would rush in and buy, thus making up for the (by assumption) "missing" demand from hedge-buyers; these arbitrageurs then close out synthetically, via spot and money markets or via other forward currency and forward money markets. So instead of a sharp price drop in one segment, we might see a tiny drop in all related markets, or even no drop at all. In fact, the hedgers themselves probably do some of the "arbitrage" work (in the wider sense), since their shopping-around calculations would normally already divert part of the selling towards spot markets if forward rates drop too deep relative to spot prices.

This role of spreading the pressure of course works for any shock, not just the forward disequilibrium we just used as an example. Suppose for instance that a central bank starts selling dollars for euros in a massive way. This would in a first instance affect the spot value: market makers see a constant flow of sell orders coming, which clogs up their books—so they lower their quotes to discourage the seller(s?) and attract new buyers. But, at constant interest rates, all forward rates would start moving also, thus also similarly influencing players in forward markets too: there is less supply, and more demand, for these slightly cheaper forward dollars. The pressure can even be born by other currencies too. For instance, suppose the market sees the change in the usd/euro rate as a dollar problem; that is, they see no good reason why the euro/yen rate would change, for instance, or the euro/gbp rate etc. Part of the pressure is diverted to yen and pound spot markets and thence to all yen and pound forward markets too, and so on. Spreading pressure helps to dampen the impact the initial spot sales wave would have had if there had been an isolated market.

The above looks at the markets as a self-centered system where hedgers place orders for exogenous reasons and where market makers just react to order flow. The role of speculators, then, is to link prices to the rest of the world. Notably, the forward price is also a risk-adjusted expected future value. So when the forward dollar depreciates while investors see no good arguments why it should, they would start buying forward, thus limiting the deviation between the forward value and the expected future value.[20] Again, this "speculative" function is a role that can be assumed by a "hedger" too; for instance, if the forward is already pretty low relative to expectations, potential hedgers of long positions may get second thoughts and decide not to sell forward after all, while players with short positions would see the extra expected gain as a nice boon that might tilt the balance in favor of hedging.

If hedgers also function as arbitrageurs (when shopping around) or as speculators (when judging the expected cost of closing out), does that mean that the usual trichotomy of players is misleading? Well, hedgers *are* special, or distinct: they start from a long or short position that has been dictated by others, like the sales or procurement departments, and they have to deal with this optimally. Speculators do not have such an exogenous motivation. But both will look at expected deviations between forward prices and expected future spot rates—"speculation"—and both will do their trades in the most economical way, thus spreading shocks into related markets—"arbitrage". So speculation and arbitrage are roles, or functions, that should be assumed by all sapient humans, including hedgers.

We are now ready to move to two related instruments—younger cousins, in fact, to forward contracts: futures, and swaps.

[20]If they take big positions, then they also assume more risk, so the risk correction may go up, too. This explains why, even at constant expectations, the forward rate may move. The point is that the discrepancy should be limited, though.

5.8 Test Your Understanding

5.8.1 Quiz Questions

1. Which of the following are risks that arise when you hedge by buying a forward contract in imperfect financial markets?

   (a) Credit risk: the risk that the counterpart to a forward contract defaults.

   (b) Hedging risk: the risk that you are not able to find a counterpart for your forward contract if you want to close out early.

   (c) Reverse risk: the risk that results from a sudden unhedged position because the counterpart to your forward contract defaults.

   (d) Spot rate risk: the risk that the spot rate has changed once you have signed a forward contract.

2. Which of the following statements are true?

   (a) Margin is a payment to the bank to compensate it for taking on credit risk.

   (b) If you hold a forward purchase contract for JPY that you wish to reverse, and the JPY has increased in value, you owe the bank the discounted difference between the current forward rate and the historic forward rate, that is, the market value.

   (c) If the balance in your margin account is not sufficient to cover the losses on your forward contract and you fail to post additional margin, the bank must speculate in order to recover the losses.

3. Which of the following statements are correct?

   (a) A forward purchase contract can be replicated by: borrowing foreign currency, converting it to domestic currency, and investing the domestic currency.

   (b) A forward purchase contract can be replicated by: borrowing domestic currency, converting it to foreign currency, and investing the foreign currency.

   (c) A forward sale contract can be replicated by: borrowing foreign currency, converting it to domestic currency, and investing the domestic currency.

   (d) A forward sale contract can be replicated by: borrowing domestic currency, converting it to foreign currency, and investing the foreign currency.

4. The following spot and forward rates are in units of THB/FC. The forward spread is quoted in centimes.

	Spot	1 month		3 month		6 month		12 month	
1 BRL	18.20–18.30	+0.6	+0.8	+2.1	+2.7	+3.8	+4.9	+6.9	+9.1
1 DKK	5.95–6.01	–0.1	–0.2	–0.3	–0.1	–0.7	–0.3	–0.9	+0.1
1 CHF	24.08–24.24	+3.3	+3.7	+9.9	+10.8	+19.3	+21.1	+36.2	+39.7
100 JPY	33.38–33.52	+9.5	+9.9	+28.9	+30.0	+55.2	+57.5	+99.0	+105.0
1 EUR	39.56–39.79	–1.7	–1.0	–3.4	–1.8	–5.8	–2.9	–10.5	–5.2

Choose the correct answer.

i. The one-month forward bid/ask quotes for CHF are:

a. 27.387–27.942 b. 25.078–24.357 c. 24.113–24.277 d. 24.410–24.610

ii. The three-month forward bid/ask quotes for EUR are:

a. 39.526–39.772 b. 36.167–37.992 c. 39.641–40.158 d. 39.397–39.699

iii. The six-month forward bid/ask quotes for JPY are:

a. 38.902–39.273 b. 88.584–91.025 c. 33.686–33.827 d. 33.932–34.095

iv. The twelve-month forward bid/ask quotes for BRL are:

a. 18.731–19.352 b. 25.113–27.404 c. 17.305–17.716 d. 18.279–18.391

5. Suppose that you are quoted the following NZD/FC spot and forward rates:

	Spot bid-ask	3-mo. forward bid-ask	*p.a.* 3 month Euro-interest	6-mo.forward bid-ask	*p.a.* 6 month Euro-interest
NZD			5.65–5.90		5.47-5.82
USD	0.5791–0.5835	0.5821–0.5867	3.63–3.88	0.5839–0.5895	3.94–4.19
EUR	0.5120–0.5159	0.5103–0.5142	6.08–6.33	0.5101–0.5146	5.60–6.25
DKK	3.3890–3.4150	3.3350–3.4410	6.05–6.30	3.3720–3.4110	5.93–6.18
CAD	0.5973–0.6033	0.5987–0.6025	1.71–1.96	0.5023–0.5099	2.47–2.75
GBP	0.3924–0.3954	0.3933–0.3989	5.09–5.34	0.3929–0.3001	5.10–5.35

(a) What are the three-month synthetic-forward NZD/USD bid-ask rates?

(b) What are the six-month synthetic-forward NZD/EUR bid-ask rates?

(c) What are the six-month synthetic-forward NZD/DKK bid-ask rates?

(d) What are the three-month synthetic-forward NZD/CAD bid-ask rates?

(e) In a–d, are there any arbitrage opportunities? What about least cost dealing at the synthetic rate?

6. True or False: Occasionally arbitrage bounds are violated using domestic ("on-shore") interest rates because:

(a) Offshore or euromarkets are perfect markets while "on-shore" markets are imperfect.

(b) Offshore or euromarkets are efficient markets while "on-shore" markets are inefficient.

5.8.2 Applications

1. Michael Milkem, an ambitious MBA student from Anchorage, Alaska, is looking for free lunches on the foreign exchange markets. Keeping his eyes glued to his Reuters screen until the wee hours, he spots the following quotes in Tokyo:

Exchange rate: Spot	NZD/USD	1.59–1.60	JPY/USD	100–101
	NZD/GBP	2.25–2.26	JPY/GBP	150–152
180 - day Forward	NZD/USD	1.615–1.626	JPY/USD	97.96–98.42
	NZD/GBP	2.265–2.274	JPY/GBP	146.93–149.19
Interest rates (simple, *p.a.*)				
180 days	USD	5%–5.25%	JPY	3%–3.25%
	NZD	8%–8.25%	GBP	7%–7.25%

Given the above quotes, can Michael find any arbitrage opportunities?

2. US-based Polyglot Industries will send its employee Jack Pundit to study Danish in an intensive training course in Copenhagen. Jack will need DKK 10,000 at t = 3 months when classes begin, and DKK 6,000 at t = 6 months, t = 9 months, and t = 12 months to cover his tuition and living expenses. The exchange rates and p.a. interest rates are as follows:

DKK/USD	Exchange rate	*p.a.* interest rate USD	*p.a.* interest rate DKK
Spot	5.820–5.830		
90 days	5.765–5.770	3.82–4.07	8.09–8.35
180 days	5.713–5.720	3.94–4.19	8.00–8.26
270 days	5.660–5.680	4.13–4.38	7.99–8.24
360 days	5.640–5.670	4.50–4.75	7.83–8.09

Polyglot wants to lock in the DKK value of Jack's expenses. Is the company indifferent between buying DKK forward and investing in DKK for each time period that he should receive his allowance?

3. Check analytically that a money-market hedge replicates an outright forward transaction. Analyze, for instance, a forward sale of DKK 1 against NZD.

Exercises 4 through 6 use the following time-0 data for the fictitious currency, the Walloon Franc (WAF) and the Flemish Yen (FLY), on Jan. 1, 2000. The spot exchange rate is 1 WAF/FLY.

	Interest rates		Swap rate
	FLY	WAF	WAF/FLY
180 days	5%	10.125%	0.025
360 days	5%	10.250%	0.050

4. On June 1, 2000, the FLY has depreciated to WAF 0.90, but the six-month interest rates have not changed. In early 2001, the FLY is back at par. Compute the gain or loss (and the cumulative gain or loss) on two consecutive 180-day forward sales (the first one is signed on Jan. 1, 2000), when you start with a FLY 500,000 forward sale. First do the computations without increasing the size of the forward contract. Then verify how the results are affected if you do increase the contract size, at the roll-over date, by a factor $1 + r^*_{T_1,T_2}$—that is, from FLY 500,000 to FLY 512,500.

5. Repeat the previous exercise, except that after six months the exchange rate is at WAF/FLY 1, not 0.9.

6. Compare the analyses in Exercises 4 and 5 with a rolled-over money-market hedge. That is, what would have been the result if you had borrowed WAF for six months (with conversion and investment of FLY—the money-market replication of a six-month forward sale), and then rolled-over (that is, renewed) the WAF loan and the FLY deposit, principal plus interest?

Chapter 6

The Market for Currency Futures

In Part 1 we first studied Interest Rate Parity (or Covered Interest Parity) in perfect markets, but we soon introduced transaction costs and other market imperfections that make life more exciting. But spreads, taxes, and information costs are not the only practical issues that can arise in this context. In this chapter, we start from two other problems connected to forwards: default risk, and absence of a secondary market. We discuss how they are handled (or not handled) in forward markets—traditionally by rationing and up-front collateralizing, nowadays also by periodic recontracting or variable collateralizing. This is the material for Section 3.

We then describe—Section 4—the institutional aspects of futures contracts. A crucial feature is that futures contracts address the problem of default risk in a way of their own: daily marking to market. This is similar to daily recontracting of a forward contract, except that the undiscounted change in the futures price is paid out in cash. In Section 5, we then trace the implications of daily marking to market for futures prices. Especially, we show that the interim cash flows from marking to market create interest risk, which affects the futures prices. In Section 6 we address the question how to hedge with futures contracts. We conclude, in the seventh section, by describing the advantages and disadvantages of using futures compared to forward contracts. In the appendix we digress on interest-rate futures—not strictly an international-finance contract, but one that is close to the FRA's discussed in Chapter 4 which, you will remember, are very related to currency forwards and forward-forward swaps.

6.1 Handling Default Risk in Forward Markets: Old & New Tricks

Futures contracts are designed to minimize the problems arising from default risk and to facilitate liquidity in secondary dealing. The best way to understand these contracts is to compare them with forward markets, where these problems also arise. When asked for forward contracts, bankers of course do worry about default by their customers—and, as we shall see, the credit-risk problem also makes it difficult to organize a secondary market for standard forward contracts. The old ways to handle default risk are rationing (refusing shady customers, that is) and asking for up-front collateral. More recent techniques are periodic re-contracting and variable collaterizing.

6.1.1 Default Risk and Illiquidity of Forward Contracts

As we saw in Part 1, a forward contract has two "legs": on the maturity date of the contract, the bank promises to pay a known amount of one currency, and the customer promises to pay a known amount of another currency. Each of these legs can be replicated by a money-market position—at least in terms of promises, that is, or as long as there is no default. However, it is important to understand that, from a bank's point of view, the credit risk present in a forward contract is of a different nature than the credit risk present in a loan. Specifically, the implicit loan and the deposit are tied to each other by the *right of offset.* The right of offset allows the bank to withhold its promised payment without being in breach of contract, should the customer default. That is, if the customer fails to deliver foreign currency (worth $\tilde{S}_T$), the bank can withhold its promised payment $F_{t_0,T}$. The bank's net opportunity loss then is $\tilde{S}_T - F_{t_0,T}$, not $\tilde{S}_T$. Likewise, if a customer bought forward but fails to pay, the bank refuses to deliver and instead sells the currency spot to the first comer; so what is at stake is again the difference between the price obtained in the cash market ($\tilde{S}_T$) and the one originally promised by the customer ($F_{t_0,T}$).

Example 6.1

Company C bought forward USD 1m against EUR. The bank, which has to deliver USD 1m, bought that amount in the interbank market to hedge its position. If Company C defaults, the bank has the right to withhold the delivery of the USD 1m. However, the bank still has to take delivery of (and pay for) the USD it had agreed to buy in the interbank market at a price $F_{t_0,T}$. Having received the (now unwanted) USD, the bank has no choice but to sell these USD in the spot market. Given default by C, the bank therefore has a risky cash flow of ($\tilde{S}_T - F_{t_0,T}$).

The second problem with forward contracts is the lack of secondary markets. Suppose you wish to get rid of an outstanding forward contract. For instance, you have a customer who promised to pay you foreign currency three months from now and, accordingly, you sold forward the foreign currency revenue to hedge the

A/R. Now you discover that your customer is bankrupt. In such a situation, you probably do not want to hold the outstanding hedge contract for another three months because the default has turned this forward position from a hedge into an open ("speculative") position. So you probably want to liquidate your forward position. Similarly, a speculator would often like to terminate a previous engagement before it matures, whether to cut her losses or to lock in her gains.

Whatever your motive of getting out early, "selling" the original forward contract is difficult. There is no organized market where you can auction off your contract: rather, you have to go beg your banker to agree on an early settlement in cash. One reason why there is no organized market is that each contract is tailor-made in terms of its maturity and contract amount, and not many people are likely to be interested in specifically your contract. Also, for your contract you probably had to provide extra security to cover default risk (see below). This means that your bank may not want you to be replaced by somebody else as a counterpart, unless comparable security is arranged (a hassle!) or you yourself guarantee the payment (dangerous!). Thus, the problem of illiquidity is partly explained by the credit-risk problem.

Example 6.2

Suppose a Spanish wine merchant received an order for ten casks of 1938 Amontillado, worth USD 1,234,567.89 and payable in 90 days, from a (then) rich American, Don Bump. The Spanish merchant hedged this transaction by selling the USD forward. However, after 35 days, Don Bump goes bankrupt (again) and will obviously be unable to pay for the wine. The exporter would like to get out of the forward contract, but it is not easy to find someone else who also wants to sell forward exactly USD 1,234,567.89 for 55 days from the current date. In addition, the wine merchant would have to convince his banker that the new counterparty is at least as creditworthy as himself.

6.1.2 Standard Ways of Reducing Default Risk in the Forward Market

As you might perhaps remember from the preceding chapter, banks have come up with various solutions that partially solve the problem of default risk: the right of offset; credit lines (when dealing with banks), or credit agreements and security (when dealing with other customers); restricted applications; and shorter lives, with an option to roll over if all goes well.

From that discussion, we see that the problem of credit risk is more or less solved by restricting access to the forward market, by requiring margins and pledges, and by limiting the maturities of forward contracts. But the second problem—illiquidity arising from the absence of secondary markets—is not addressed. One can negotiate an early (premature) settlement with the original counterparty of the forward contract. But this is a question of negotiation, not a built-in right for the holder of the contract. Also, one cannot rely on an immediately observable market price to determine the value of the outstanding contract. Rather, one has to compute the bounds on the fair value (using the Law of the Worst Possible Combination), and negotiate some price within these bounds. Thus, the early settlement of forward contracts is rather inconvenient. As a result, and in contrast to futures contracts, virtually all forward contracts remain outstanding until they expire, and actual delivery and payment is the rule rather than the exception. Closing out, if done at all, often is via adding a reverse contract, as we have seen. While this works out well enough most of the time, a long and a short do not add up to a zero position if there is default:

Example 6.3

Some time ago you bought USD 15m forward from the Herstatt & Franklin, your favorite bank, but you have just closed out by selling to it, same amount and same date. You think you're out; however, if prior to T H&F have gone into receivership, then you have a problem. One of the two contracts probably has a negative value to you and the other a positive one. Then the bank's receivers will make you pay for the one with the negative value. For the contract with a positive value, though, you can only file a claim with the receivers, and *maybe* you'll see part of your money some day.

6.1.3 Reducing Default Risk by Variable Collateral or Periodic Recontracting

As we saw in the previous section, one often needs to post margin when a forward contract is bought or sold. The margin may consist of an interest-earning term deposit, or of securities (like stock or bonds). Please note that posting margin is very different from paying something to the bank. A payment is made to settle a debt, or to become the owner of a commodity or a financial asset. Whatever the reason for the payment, the bank that receives a payment becomes the owner of the money. In contrast, margin that is posted still belongs to the customer; the bank or broker merely has the right to seize the collateral if and only if the customer defaults.

The required margin can be quite high because the bank is willing to take only a small chance that the contract's expiration value, if negative, is not covered by the margin. In about half of the cases, the collateral will turn out to have been unnecessary because there is roughly a 50 percent chance that $\tilde{S}_T - F_{t_0,T}$ will end up being positive. There are two ways to reduce the need for margin.

- **Variable collateral** Under this system, the bank requests two kinds of margin. First, there is a small but permanent margin—say, the amount that almost surely covers the worst possible one-day drop in the market value of the forward contract. If the market value of the contract becomes negative, the bank then asks for additional collateral in order to cover at least the drop in the current market value of the forward contract. If the customer fails to put up the additional margin, the bank seizes all margin put up in the past—including the initial safety margin—and closes out the outstanding contract in the forward market. Obviously, under such a system the amount of collateral that has to be put up is far smaller, on average, than what is required if a single, large initial margin has to be posted. The reason is that, under this system, collateral is called for only when needed, and only to the extent that it is needed at that time.

- **Periodic recontracting**. Under this system, the new market value of yesterday's contract is computed every day. The party that ends up with a negative value then buys back the contract from the counterparty, and both sign a new contract at the day's new price. If the loser fails to settle the value of yesterday's contract, the bank seizes the initial margin, and closes out the contract in the forward market. Under this system, only a small amount of margin is needed, since the collateral has to cover only a one-day change in the market value.

It is useful to spell out the cash flows, because this will help you understand what futures contracts are and why they differ from recontracted forwards.

Example 6.4

Suppose that, at time 0, Smitha Steel has bought forward USD against INR for delivery at time 3. In Table 6.1 we describe the implications under the systems of variable collateral and periodic contracting, respectively. We ignore the initial margin, since it is the same in both cases. All amounts are in INR. The example assumes that the forward rate always goes down, as this is the possibility that Smitha's bank worries about.

With variable collateral, nothing is changed relative to a standard contract (except that collateral is asked only if and when needed): Smitha has temporarily moved some assets from her own safe to a safe of her bank, but gets them back upon paying INR 40m at time 3, as promised under the forward contract. With recontracting, in contrast, there are three genuine payments, one per day, but by design their time-value-corrected final value is still equal to INR 40m at time 3. To see this, just consider the total paid, at time 3, when the interim losses are financed by loans which are paid back at time 3:

$$\begin{array}{lr}
\text{- time 1: pay } \underbrace{(40-38)/1.02}_{\text{1.961 borrowed}} \times 1.02 = 40 - 38 = & 2 \\
\text{- time 2: pay } \underbrace{(38-36)/1.01}_{\text{1.98 borrowed}} \times 1.01 = 38 - 36 = & 2 \\
\text{- time 3: pay} & \underline{36} \\
\text{- total: pay} & 40
\end{array}$$

So the discounting, which is part of the market value calculations that are behind the recontracting payments, also means that after taking into account time value the recontracting cancels out: it can be "undone" by financing any losses via loans, or by depositing any gains, thus shifting all cashflows back to time $T = 3$.

Table 6.1: **Forward Contracts with Variable Collateral or Daily Recontracting**

data	Variable Collateral	Periodic Recontracting
time 0: $F_{0,3} = 40$ $r_{0,3} = 3\%$	Smitha buys forward USD 1m at $F_{0,3} = 40$	Smitha buys forward USD 1m at $F_{0,3} = 40$
time 1: $F_{1,3} = 38$ $r_{1,3} = 2\%$	Market value of old contract is $\frac{38m-40m}{1.02} = -1.961m$ Smitha puts up T-bills worth at least 1.961m	Market value of old contract is $\frac{38m-40m}{1.02} = -1.961m$ Smitha buys back the old contract for 1.961m and signs a new contract at $F_{1,3} = 38$.
time 2: $F_{1,3} = 36$ $r_{2,3} = 1\%$	Market value of old contract is $\frac{36m-40m}{1.01} = -3.960m$ Smitha increases the T-bills put up to at least 3.960m	Market value of old contract is $\frac{36m-38m}{1.01} = -1.980m$ Smitha buys back the old contract for 1.980m and signs a new contract at $F_{2,3} = 36$.
time 3: $F_{3,3} = S_3 = 34$ $r_{3,3} = 0\%$	Smitha pays the promised INR 40m for the USD 1m, and gets back her T-bills	Smitha pays the promised INR 36m for the USD 1m
total paid:	INR 40m	(adjusted for time value:) - time 3: 36m - time 2: 1.980 × 1.01= 2m - time 1: 1.961 × 1.02= 2m - total: 40m

DoItYourself problem 6.1

Given a sequence $\{F_{1,4},\ F_{2,4},\ F_{3,4},\ F_{4,4} = S_4\}$, write in algebra the cash flows from daily recontracting, and show that if all losses are financed by loans and all gains are deposited until time $T = 4$, you pay, all in all, $F_{1,4}$.

The system of variable collateral is used in many stock exchanges in continental Europe. Somewhat confusingly, these contracts are sometimes called futures contracts; in reality, they are collateralized forward contracts. “Futures” just sounds cooler than forwards, though.

This finishes our discussion of credit risks in forward contracts. We now see how this is handled in futures markets, and how secondary dealing has been organized.

6.2 How Futures Contracts Differ from Forward Markets

A currency futures contract has the following key characteristics: (*i*) it has zero initial value; (*ii*) it stipulates delivery of a known number of forex units on a known future date T; and (*iii*) the HC payment for the forex is a known amount $f_{t,T}$, paid later.

The only news here, relative to a forward contract, is the last word—the vague term "later" rather than the precise expression "at T". In fact, we can be more specific about the timing of the payments: of the total, which is $f_{t,T}$, the part $f_{t,T} - \tilde{S}_T$ is paid gradually during the life of the contract via daily marking-to-market payments, and the remainder, $\tilde{S}_T$, is paid at maturity. Note that the pattern of the payments over time is ex ante unknown: we only know the grand total that we will pay, the no-time-value-correction sum.

We show how this marking-to-market system is a somewhat primitive version of the daily-recontracting system we discussed in the previous section. So it is a way to mitigate the problem of default risk. Given that this problem is largely solved, futures contracts can be transferred among investors with minimal problems. We'll see how this is done: with standardized contracts, in organized markets, and with the clearing corporation as the central counterpart. We'll use the following jargon: "buying a contract" means engaging in a purchase transaction—going long forex, that is: you will get forex and pay HC; and a futures price is per unit of currency, even though the contract always is for a multiple of FC units.

6.2.1 Marking to Market

Recall that when a forward contract is recontracted every day, the buyer receives a daily cash flow equal to the *discounted* change in the forward price. Thus, rising prices mean cash inflows for the buyer, and falling prices mean cash outflows. (The signs are reversed when the seller's point of view is taken.) Also, as the interim payments are based on the discounted forward price, the total amount paid is still equivalent to paying the initially contracted rate, $F_{t_0,T}$, at the contract's expiration date.

A futures contract works quite similarly, except that the discounting is omitted. So the daily payments are equal to the undiscounted changes in the futures prices. The reasons for this simplifications were not hard to guess: it made sense at the time futures were designed, the mid 1800s. (*i*) Futures contracts had short lives, and interest rates were low (these were the days of the gold standard), so discounting made no huge difference. (*ii*) Discounting means smaller payments; this is welcome when the payment is an outflow (like in our Smitha example), but it's bad news when we face inflows. So if price rises are roughly equally probable as price falls, on average it made no difference, people felt. And (*iii*), painfully, in the 1800s discounting would have to be done manually rather than electronically. For these reasons people simply dropped it. As we shall see, the argument that "it all washes out as price rises are as probable as falls" is not quite true, but the effect is indeed minimal.

So in practice we have daily cash flows that, for the buyer, are equal to $f_{t,T} - f_{t-1,T}$, with the final payment, $f_{T,T} = S_T$, taking place after the last trading day. The last trading day is two working days before delivery, like in spot markets. So the last-trading-day futures price must be equal to the contemporaneous spot rate. As a result, after all the marking-to-market payment have been made, the buyer is left with a spot contract.

Example 6.5

In the Smitha Steel example, suppose the rates were futures prices rather than forward ones. Then the cash flows would have been -2, -2, -2 (= the last marking to market), and -34 (the spot payment, $f_{3,3} = S_3$). Below, I detail this, and compare it to a periodically recontracted forward contract:

price rate r	40; r=0.03	38 r=0.02	36 r=0.01	34 r=0.00
futures	–	$38 - 40 = -2$	$36 - 38 = -2$	$34 - 36 = -2$ and then buy at 34
recontracted fwd	–	$\frac{38-40}{1.02} = -1.961$	$\frac{36-38}{1.01} = -1.980$	buy at 36

Thus, ignoring time value, the cumulative payments from the buyer are equal to 40 units of home currency.

The cash flows to the seller are the reverse. In fact, what happens is that the buyer pays the seller if prices go down and receives money from the seller if prices go up. In short, good news (rising prices for the buyer, falling prices for the seller) means an immediate inflow, and bad news an immediate outflow. These daily payments from "winner" to "loser" occur through accounts the customers hold with their brokers, and they are transmitted from the loser to the winner through brokers, *clearing members*, and the clearing corporation. The *settlement price*, upon which the daily marking-to-market cash flows are based, is in principle equal to the day's closing price or close price. However, futures exchanges want to make sure that the settlement price is not manipulated; or they may want a more up-to-date price if the last transaction took place too long before the close. One way to ensure this is to base the settlement price not on the actual last trade price but on the average of the transaction prices in the last half hour of trading or, if there is no trading, the average of the market makers' quotes (LIFFE).

Suppose, lastly, that somewhere in the middle of the second trading day, the day where the price drops from 38 to 36, Smitha sells her contract at a forward price 37.5. The total marking to market for day 2 is still 36–38=–2; but now this will be split into 37.5–38=–0.5 for Smitha, and 36–37.5 = –1.5 for the (then unsuspecting) new holder.

Marking to market is the most crucial difference between forward and futures contracts. It means that if an investor defaults, the "gain" from defaulting is simply the avoidance of a one-day marking-to-market outflow: all previous losses have already been settled in cash. This implies the following:

- Compared to a forward contract, the incentive to default on a futures contract is smaller. By defaulting on the marking-to-market payment, one avoids only a payment equal to that day's price change. In contrast, in the case of a forward contract, defaulting means that the investor saves the amount lost over the entire life of the contract.

  Example 6.6

  Investor A bought EUR 1m at $f_{t_0,T}$ = USD/EUR 0.96. By the last day of trading but one, the futures price has drifted down to a level of USD/EUR 0.89. So investor A has already paid, cumulatively, 1m × (0.96 – 0.89) = USD 70,000 as marking-to-market cash flows. If, on the last day of trading, the price moves down by another ten points, then, by defaulting, investor A avoids only the additional payment of 1m × 0.001 = USD 1,000. In contrast, if this had been a forward contract, the savings from defaulting would have been the entire price drop between t_0 and T, that is, 1m × $(F_{t_0,T} - S_T)$ = 1m × (0.96 – 0.889) = USD 71,000.

- From the point of view of the clearing house, the counterpart of the above statement is that if an investor nevertheless fails to make the required margin payment, the loss to the clearing house is simply the day's price change.

In practice, the savings from defaulting on a futures contract (and the clearing house's loss if there is default) are even smaller than the above statement suggests because of a second characteristic of futures markets—the margin requirements.

6.2.2 Margin Requirements

To reduce even the incentive of evading today's losses, the buyer or seller also has to put up initial security that almost surely covers a one-day loss. This is true security in the sense that one earns interest on it.[1] The general idea behind the margin requirements is that the margin paid should cover virtually all of the one-day risk. This, of course, further reduces both one's incentive to default as well as the loss to the clearing house if there is default.

Margin also means limit, or line. In that sense, two margins have to be watched when trading in futures markets, *initial margin* and *maintenance margin.* Indeed, in

[1] The marking-to-market payments are often called margin payments. This term is a bit misleading if "paying margin" is interpreted as "posting additional security": if the payments really were security, the payer would still be entitled to the normal interest on the money put up. In reality there are no interest payments on the m-to-m payments, so economically these are final payments not security postings—unlike the initial margin, which is genuine security.

theory every gain or loss is immediately settled in cash, but this may mean frequent, small payments which is costly and inconvenient. So in practice losses are allowed to accumulate to certain levels before a *margin call* (a request for payment) is issued. These small losses are simply deducted from the initial margin until a lower bound, the maintenance margin, is reached. At this stage, a margin call is issued, requesting the investor to bring the margin back up to the initial level.

Example 6.7

The initial margin on a GBP 62,500 contract may be USD 3,000, and the maintenance margin USD 2,400. The initial USD 3,000 margin is the initial *equity* in your account. The buyer's equity increases (decreases) when prices rise (fall), that is, when marking-to-market gains or losses are credited or debited to your account. As long as the investor's losses do not sum to more than USD 600 (that is, as long as the investor's equity does not fall below the maintenance margin, USD 2,400), no margin call will be issued to her. If her equity, however, falls below USD 2,400, she must immediately add *variation margin* to restore her equity to USD 3,000.

Failure to make the margin payment is interpreted as an order to liquidate the position. That is, if you bought and cannot pay, your contract will be put up for sale at the next opening, as if you had ordered to sell the contract; and if you were short, your contract will likewise be closed out the next day as if you had ordered to buy. This way, the Exchange finds a new party that steps into your shoes. The loss or gain on this last deal is yours, and is added to or subtracted from the margin.

Example 6.8

When Nick Leeson had gambled his employer, the then 233-year-old Barings bank, into ruin he had accumulated losses of GBP 800m, more than Barings' entire equity. But the Singapore Exchange lost "only" 50m. Barings London had sent Nick about 500m for m-to-m payments (thinking these were deposits or something like that), and Nick had "borrowed" about 250m from other customers' accounts to pay even more margin without telling London. So the SME was already covered for about GBP 750m. The balance was lost when Nick's huge open positions were liquidated at short notice and when the initial margin proved totally inadequate to cover the losses caused by the massive price pressure.

6.2.3 Organized Markets

As we saw, forward contracts are not really traded; they are simply initiated in the over-the-counter market (typically with the client's bank) and held until maturity. In the forward market, market makers quote prices but there is no organized way of centralizing demand and supply. The only mechanisms that tend to equalize the prices quoted by different market makers are arbitrage and least cost dealing; and, as traders are in permanent contact with only a few market makers, price equalization

Leeson's Lessons *re* Barings' End

What went wrong in the Barings case, and can it happen again? Both the futures exchanges and Barings (and possibly many other firms, in those days) made a number of mistakes:

- **Internal organisational problems** Nick Leeson (above) headed both the dealing room (front office) and the accounting interface (back office). Also, he came from the back office. So he could bend the rules, key in misleading records, and funnel cash between various accounts. Also, there was no middle office (risk management) and there were no enforced position limits.
- **Gullible greed in London** Barings' HQ thought Nick was making huge profits and did not want to slaughter the goose with the golden eggs, so they kept sending money which they thought was just security postings.
- **Failing oversight** Both the Osaka and Singapore futures exchanges were worried about the size of Nick's positions, and they talked about it to each other, but in the end did nothing.

These mistakes are unlikely to be made again any time soon in any well-run firm. That is, the next catastrophe will again be of a totally unexpected nature.

That, at least, is what we all thought. Yet in January 2008 it transpired that Jerôme Kerviel at France's *Société Générale* had built a secret portfolio of stock futures for a notional value of EUR 50b—more than the bank's own market value of equity then, 36b—on which the realized loss turned out to be 4.9b. Of the total loss, his lawyers objected, almost two thirds was due to a panic liquidation by SG after discovering a 1.7b proper loss by Kerviel himself.

Before being a trader he had worked in IT in the middle office, where he had figured out five passwords and identified some loopholes. For instance, SG checked the position limits every three days only; so just before the checks, Kerviel simply reduced the net exposures by ficticious trades. (Checks should be random and frequent, and limits should look not just at net but also gross positions.) Worse, SG was blamed for ignoring no fewer than 75 danger signals (including 'does not take up his holidays' and 'sweats a lot', alongside, more seriously, worried questions from futures exchanges). Like Barings before, SG preferred to look the other way because Kerviel had posted a profit of 1.4b in 2006 (on a maximum position of 125m!?).

is imperfect. Nor is there any public information about when a transaction took place, and at what price.

In contrast, futures are traded on organized exchanges, with specific rules about the terms of the contracts, and with an active secondary market. Futures prices are the result of a centralized, organized matching of demand and supply. One method of organizing this matching of orders is the *open outcry* system, where floor members are physically present in a trading pit and auction off their orders by shouting them out. US exchanges traditionally work like this; so did London's LIFFE and Paris' MATIF.[2] You can see open outcry trading in an Ackroyd-Murphy movie

[2] LIFFE: London International Financial Futures Exchange (where also options are traded, since

Trading Places. Another method, traditionally used in some continental exchanges (including Germany's DTB and Belgium's Belfox, now part of Eurex and Euronext, respectively) is to centralize the limit orders in a computerized *Public Limit Order Book.*[3] Brokers sit before their screens, and can add or delete their orders, or fill a limit order posted on the screen. Computerized trading, whether *price-driven* (i.e. with market makers) or *order-driven* (with a limit-order book) is gradually replacing the chaotic, intransparant open-outcry system.

6.2.4 Standardized Contracts

Each forward contract is unique in terms of size, and the expiry date can be chosen freely. This is convenient for hedgers that mean to hold the contract until maturity, but unhandy if secondary markets are to be organized: for every single trade, new terms and conditions would have to be keyed in and new interest rates dug up.

To facilitate secondary trading, all futures contracts are standardized by contract size (see Table 6.2 for some examples) and expiration dates. This means that the futures market is not as fragmented—by too wide a variety of expiration dates and contract sizes—as the forward market. Although standardization in itself does not guarantee a high volume, it does facilitate the emergence of a deep, liquid market.

Expiration dates traditionally were the third Wednesdays of March, June, September, or December, or the first business day after such a Wednesday. Nowadays, longer-lived contracts and—for the nearer dates—a wider range of expiry dates are offered, but most of the interest still is for the shortest-lived contracts. Actual delivery takes place on the second business day after the expiration date. When a contract has come to expiry, trade in a distant-date contract is added. For instance, in the old March-June-Sept-Dec cycle, the year starts with March, June, and September contracts, but as soon as the March contract is over one adds a December contract to the menu, and so on.

---

the merger with the London Traded Options Exchange). MATIF: *Marché à Terme International de France* (International Futures Market of France). Both are now part of Euronext, which has grouped all its futures and options business in London, under Euronext LIFFE.

[3]DTB: Deutsche Termin Börse. Belfox: Belgian Futures and Options Exchange. A limit order is an order to buy an indicated number of currency units at a price no higher than a given level, or to sell an indicated number of currency units at a price no lower than a given level. The limit orders submitted by an individual reveal the individual's supply and demand curve for the currency. By aggregating all limit orders across investors, the market supply and demand curves are obtained. The market opens with a call, that is, with a computer-determined price that equates demand and supply as closely as possible. Afterwards, the computer screens display the first few unfilled orders on each side (purchase orders, and sell orders), and brokers can respond to these, or cancel their own orders, or add new orders as customer orders come in.

Table 6.2: **Contract sizes at some futures exchanges**

Rate	at	contract size (FC)	Other exchanges
USD/GBP	IMM	62,500	PBOT, LIFFE, SIMEX, MACE
USD/EUR	IMM	125,000	LIFFE, PBOT, SIMEX, MACE, FINEX
EUR/USD	OM-S	50,000	EUREX
USD/CHF	IMM	125,000	LIFFE, MACE, PBOT
USD/AUD	IMM	100,000	PBOT, EUREX
NZD/USD	NZFE	50,000	
USD/NZD	NZFE	100,000	
USD/JPY	IMM	12500,000	LIFFE, TIFFE, MACE, PBOT, SIMEX
USD/CAD	IMM	100,000	PBOT, MACE

Key
EUREX = European Exchange (comprising the former German DTB and the former Swiss SFX)
IMM-International Money Market (Merc, Chicago)
LIFFE = London International Financial Futures Exchange
MACE = MidAmerican Commodity Exchange
NZFE = New Zealand Futures Exchange.
OM-S = OptionsMarkned Stockholm
PBOT = Philadelphia Board of Trade
SIMEX = Singapore International Money Exchange
TCBOT = Twin Cities Board of Trade (St Paul / Minneapolis)
TIFFE = Tokyo International Financial Futures Exchange

Source Data from Sercu and Uppal, R., International Financial Markets and the Firm, South-Western Publishing Company, 1995.

6.2.5 The Clearing Corporation

The clearing corporation serves two purposes: acting as central counterparty (CCP)—'novation'—and clearing of an investor's offsetting trades.

Novation Formally, futures contracts are not initiated directly between individuals (or corporations) A and B. Rather, each party has a contract with the futures clearing corporation or clearing house that act as CCP. For instance, a sale from A to B is structured as a sale by A to the CCP, and then a sale by the CCP to B. Thus, even if B defaults, A is not concerned (unless the clearing house also goes bankrupt). The clearing corporation levies a small tax on all transactions, and thus has reserves that should cover losses from default.

Clearing The clearing house thus guarantees payment or delivery. In addition, it effectively "clears" offsetting trades: if A buys from B and then some time later sells to C, the clearing house cancels out both of A's contracts, and only the Clearing House's contracts with B and C remain outstanding. Player A is effectively exonerated of all obligations. In contrast, as we saw, a forward purchase by A from B and a forward sale by A to C remain separate contracts that are not cleared: if B fails to deliver to A, A has to suffer the loss and cannot invoke B's default to escape its (A's) obligations to C.

Figure 6.1: **WSJ information on currency futures**

FUTURES PRICES [...] **CURRENCY**

	Open	High	Low	Settle	Change	Lifetime High	Lifetime Low	Open Interest
JAPAN YEN (CME) — 12.5 million yen ; $ per yen (.00)								
Sept	.9458	.9466	.9386	.9389	– .0046	.9540	.7945	73,221
Dec	.9425	.9470	.9393	.9396	– .0049	.9529	.7970	3,455
Mr94				.9417	– .0051	.9490	.8700	318

Est vol 28,844; vol Wed 36,595; open int 77,028, + 1.820

6.2.6 How Futures Prices Are Reported

Figure 6.1 contains an excerpt from The Wall Street Journal, showing information on yen futures trading at the International Money Market (IMM) of the Chicago Mercantile Exchange (CME). The heading, JAPAN YEN, shows the size of the contract (12.5m yen) and somewhat obscurely tries to say that the prices are expressed in USD cents. The June 1993 contract had expired more than a month before, so the three contracts being traded on July 29, 1993, are the September and December 1993 contracts, and the March 1994 contracts. In each row, the first four prices relate to trading on Thursday, July 29—the price at the start of trading (open), the highest and lowest transaction price during the day, and the settlement price ("Settle"), which is representative of the transaction prices around the close.

The settlement price is the basis of marking to market. The column, "Change," contains the change of today's settlement price relative to yesterday. For instance, on Thursday, July 29, the settlement price of the September contract dropped by 0.0046 cents, implying that a holder of a purchase contract has lost 12.5m $\times$ (0.0046/100) = USD 575 per contract and that a seller has made USD 575 per contract. The next two columns show the highest and lowest prices that have been observed during the life of the contract. For the March contract, the "High-Low" range is more narrow than for the older contracts, since the March contract has been trading for little more than a month. "Open Interest" refers to the number of outstanding contracts. Notice how most of the trading is in the nearest-maturity contract. Open interest in the March '94 contract is minimal, and there has not even been any trading that day. (There are no open, high, and low data.) The settlement price for the March '94 contract has been set by the CME on the basis of bid-ask quotes.

The line below the price information gives an estimate of the volume traded that day and the previous day (Wednesday). Also shown are the total open interest across the three contracts, and the change in open interest relative to the day before.

* * *

This finishes our review of how futures differ from forwards. From a theoretical perspective, the main difference is the marking to market, or, if you wish, the

Table 6.3: **HC cash flows assuming that** $F_{0,2} = f_{0.2}$

	HC cash flow: futures		HC cash flow: forward		difference	
$F_{1,2}$	time 1	time 2	time 1	time 2	time 1	time 2
105	$105 - 100 = +5$	$(\tilde{S}_2 - 105) - \tilde{S}_2 = -105$	0	-100	+5	-5
95	$95 - 100 = -5$	$(\tilde{S}_2 - 95) - \tilde{S}_2 = -95$	0	-100	-5	+5

omission of discounting in the daily recontracting. In the next section we see whether this has an impact on the pricing and, if so, in what direction.

6.3 Effect of Marking to Market on Futures Prices

We saw that the absence of discounting in the daily recontracting has been waved aside as unimportant, *ex ante* at least, if price rises and price drops are equally unlikely. Is this a good argument? In this section we show that the claim is OK if price changes are independent of the time path of interest rates—which is not quite true, but is close enough for most purposes.

Recall that if a corporation hedges a foreign-currency inflow using a forward contract, there are no cash flows until the maturity date, T; and, at T, the money paid by the debtor is delivered to the bank in exchange for a known amount of home currency. In contrast, if hedging is done in the futures markets, there are daily cash flows. As we saw in the beginning of this chapter, interim cash flows do not affect pricing if these cash flows are equal to the discounted price change, as is the case in a forward contract that is recontracted periodically. The reason is that, with daily recontracting, one can "undo" without cost the effects of recontracting by investing all inflows until time T and by financing all outflows by a loan expiring at T. The question we now address is whether the price will be affected if we drop the discounting of the price changes—that is, if we go from forward markets to futures markets. We will develop our argument in three steps, and illustrate each step using an example. For simplicity, we assume that next period there are only two possible futures prices and that investors are risk neutral. All these simplifying assumptions can easily be relaxed without affecting the final conclusion.

Let there be three dates ($t = 0$, $t = 1$, and $t = T = 2$, the maturity date, and let the initial forward rate be $F_{0,2}$ = USD 100. Let there be only two possible time-1 forward prices, either 105 or 95, and let these be equally probable. We want to verify the conjecture that $f_{t,2} = F_{t,2}$. This is easily seen to be true at time 1: since as of that date there are no more extra m-to-m cashflows relative to forward contracts, futures and forward prices must be the same at time $T - 1$. The issue is whether that also holds for earlier dates—or *the* earlier date, in our case. The answer must be based on the difference of the cash flows between the two contracts (Table 6.3):

Table 6.4: **HC net time value effect at $t = 2$ assuming that $F_{0,2} = f_{0.2}$**

	case 1			case 2			case 3		
state	r	net time value at $t = 2$		r	net time value at $t = 2$		r	net time value at $t = 2$	
up	0	$5 \times 1.00 - 5 =$	0.00	0.10	$5 \times 1.10 - 5 =$	0.50	0	$5 \times 1.08 - 5 =$	0.40
down	0	$-5 \times 1.00 + 5 =$	0.00	0.10	$-5 \times 1.10 + 5 =$	-0.50	0	$-5 \times 1.12 + 5 =$	-0.60
E(.)			0.00			0.00			-0.10

- The buyer of the forward contract simply pays 100 at time 2. This is shown under the columns "Cash flows from forward," in Table 6.3.
- The buyer of the futures contract pays 5 or receives 5, depending on the price change at time 1. The balance is then paid at time $T = 2$, partly as the last m-to-m payment and partly as the HC leg of a spot purchase. Thus, the buyer will receive/pay the cash flows shown under the columns "Cash flows from futures"—either –5 and –95, or +5 and –105.
- The columns labeled "Difference in cash flows" show the cash flows for the futures contract relative to the cash flow of the forward contract.

We see that the futures is like a forward except that the buyer also gets a zero-interest loan of 5 in the upstate, and must make a zero-interest deposit of 5 in the downstate. Whether this zero-rate money-market operation makes a difference depends on interest rates. In table 6.4 we look at three cases: a zero interest rate in both the up and the down state, a 10-percent interest rate in both the up and the down state, and lastly an 8% rate in the up-state and a 12% one in the down-state.

- In the zero-rate case you of course do not mind receiving a zero-rate loan in the up-state, but you do not think this is valuable either: everybody can get that for free, by assumption. Nor do you mind the forced deposit at zero percent in the down-state: you can borrow the amount for free from a bank anyway. In short, the m-to-m flows do not add or destroy any value when interest rates are zero. It follows that the conjecture $F = f$ is acceptable.
- In the case with a 10% interest rate you positively love receiving a zero-rate loan: you can invest that money and earn 0.50 on it at time 2. In contrast, now you do mind the forced deposit at zero percent: you lose 0.50 interest on it. But if the up- and down-scenarios are equally probable, a risk-neutral investor still does not really mind, *ex ante*: the expected time-value effect remains zero. It follows that the conjecture $F = f$ is still acceptable when the risk-free rate is a positive constant.

- In the third case you still like the zero-rate loan, but the gain is lower: at time 2 you make just 8% on the 5, or 0.40. Likewise you still mind the forced deposit at zero percent, but now you lose 0.60 time value on it since the interest rate is higher, 12%. And if, *ex ante*, the up- and down-scenarios are equally probable, a risk-neutral agent now dislikes the zero-rate operations: the expected time value effect is now negative. It follows that the conjecture $F = f$ is no longer acceptable when the risk-free rate is higher in the down-state.

The example is quite special, but the basic logic holds under very general circumstances as it is based on a simple syllogism:

Fact 1 Unexpectedly low interest rates tend to go with rising asset prices, while unexpectedly high interest rates tend to go with falling prices.

Fact 2 To the futures buyer, rising prices are like receiving a zero-interest loan, relative to a forward contract, while falling prices mean zero-interest lending (you have to pay money to the clearing house).

Therefore the time-value game is not fair: you get the free loan when rates tend to be low, while you are forced to lend for free when rates tend to be high.

Stated differently, money received from marking to market is, more often than not, reinvested at low rates, while intermediate losses are, on average, financed at high rates. Thus, the financing or reinvestment of intermediary cash flows is not an actuarially fair game. If futures and forward prices were identical, a buyer of a futures contract would, therefore, be worse off than a buyer of a forward contract. It follows that, to induce investors to hold futures contracts, futures prices must be lower than forward prices.[4]

The above argument is irrefutable, and contradicts the gut feeling of the 1800s that discounting or not made no difference, on average. But how important is the effect? In practice, the empirical relationship between exchange rates and short-term interest rates is not very strong. Moreover, simulations by, for example, French (1983) and Cornell and Reinganum (1981) have shown that even when the interest rate is negatively correlated with the futures price, the price difference between the forward and the theoretical futures price remains very small—at least for short-term contracts on assets other than T-bills and bonds. Thus, for practical purposes, one can determine prices of futures contracts almost as if they were forward contracts.

Note that what is called the spot-forward swap rate in forward markets tends to be called the *basis* when we deal with futures:

$$\text{basis} := f_{t,T} - S_t. \tag{6.1}$$

[4]If the correlation were positive rather than negative, then marking to market would be an advantage to the buyer of a futures contract; as a result, the buyer would bid up the futures price above the forward price. Finally, if the correlation would be zero, futures and forward prices would be the same.

Since $f \approx F$, it follows that the basis is positive (*i.e.* futures prices are above spot prices) if the foreign rate of return is below the domestic one, and *vice versa.*

6.4 Hedging with Futures Contracts

In this section, we see how one can use futures to hedge a given position. Because of its low cost even for small orders, a hedger may prefer the currency futures market over the forward market. There are, however, problems that arise with hedging in the futures market.

- The contract size is fixed and is unlikely to exactly match the position to be hedged.
- The expiration dates of the futures contract rarely match those for the currency inflows/outflows that the contract is meant to hedge.
- The choice of underlying assets in the futures market is limited, and the currency one wishes to hedge may not have a futures contract.

That is, whereas in the forward market we can tailor the amount, the date, and the currency to a given exposed position, this is not always possible in the futures market. An imperfect hedge is called a *cross-hedge* when the currencies do not match, and is called a *delta-hedge* if the maturities do not match. When the mismatches arise simultaneously, we call this a *cross-and-delta hedge.*

Example 6.9

Suppose that, on January 1, a US exporter wants to hedge a SEK 9,000,000 inflow due on March 1 (=T_1). In the forward market, the exporter could simply sell that amount for March 1. In the futures market, hedging is less than perfect:

- There is no USD/SEK contract; the closest available hedge is the USD/EUR futures contract.
- The closest possible expiration date is, say, March 20 (T_2).
- The contract size is EUR 125,000. At the current spot rate of, say, SEK/EUR 9.3, this means SEK 1,162,500 per contract.

So assuming, irrealistically, a constant SEK/EUR cross rate, the hedger would have to sell eight contracts to approximately hedge the SEK 9,000,000: $8 \times 1,162,500 = 9.3$m. But the more difficult question is how to deal with the cross-rate uncertainty and the maturity mismatch. This is the topic of this section.

As we shall see, sometimes it is better to hedge with a portfolio of futures contracts written on different sources of risks rather than with only one type of futures contract. For example, theoretically there is an interest risk in both SEK and USD because the dates of hedge and exposure do not match, so one could consider taking futures positions in not just EUR currency but also in EUR and USD interest rates, and perhaps even SEK interest rates. However, in order to simplify the exposition, we first consider the case where only one type of futures contract is being used to hedge a given position.

6.4.1 The Generic Problem and its Theoretical Solution

The problems of currency mismatch and maturity mismatch mean that, at best, only an approximate hedge can be constructed when hedging with futures. The standard rule is to look for a futures position that minimizes the variance of the hedged cash flow. Initially, we shall assume the following:

- There is one unit of foreign currency e ("exposure") to be received at time T_1—for instance, one Swedish Krona is to be converted into USD, the HC.
- A futures contract is available for a "related" currency h ("hedge")—for instance, the EUR—with an expiration date T_2 ($\geq T_1$).
- The size of the futures contract is one unit of foreign currency h (for instance, one EUR).
- Contracts are infinitely divisible; that is, one can buy any fraction of the unit contract.

Items 1 and 3 are easily corrected. Item 4 means we will ignore the fixed-contract-size problem. The reason is that nothing can be done about it except finding a theoretical optimum and then rounding to the nearest integer.

Let us show the currency names as superscripts, parenthesized so as to avoid any possible confusion with exponents. Denote the number of contracts sold by β.[5] The total cash flow generated by the futures contracts between times t and T_1 is then given by the size of the position, $-\beta$, times the change in the futures price between times t and T_1. (True, this ignores time-value effects, but we can't be too choosy: the hedge is approximate anyway.)

Example 6.10

Boston Summae Cornucopiae Ltd will receive SEK in May, so they want to take a position in the June EUR contract to hedge this. Since a crown is worth about 10 eurocent, they could hedge each crow by 0.10 euros, meaning that β is set equal to 0.10.

If the EUR/SEK rate remains constant regardless of the gyrations of the USD/EUR rate, then each pip change in the USD/SEK rate is associated with a one-tenth-pip

[5] Beta should get a double time subscript, as should the variance and covariance in the solution. But the notation is already cluttered enough.

change in the USD/EUR rate. To hedge this one-tenth-pip change in the USD/EUR rate, one tenth of a euro then suffices. For instance, if the spot rate changes from USD/SEK 0.14 to 0.15 and the EUR/SEK rate remains at 0.10, then the euro goes up from USD/EUR 1.4 to 1.5; so holding 0.10 euro would be enough to hedge the crown position.

The above example shows you one simple rule for chosing the hedge ratio: set it equal to the relative value, the EUR/SEK cross rate in our case. The example also makes it clear that this rule assumes a fixed cross rate, which cannot be literally true. So our issue is under what assumptions the above simple rule still works and how we can do better if the rule of thumb fails. In general, the hedged cash flow equals

$$\text{Cash flow at time } T_1 \;= \tilde{S}^{(e)}_{T_1} - \beta \times (\tilde{f}^{(h)}_{T_1,T_2} - f^{(h)}_{t,T_2}). \tag{6.2}$$

Example 6.11

If Boston SC has set beta equal to 0.10, and the SEK then appreciates from 0.14 to 0.15 while the EUR appreciates from 1.40 to 1.48, Boston SC's cash flow is

$$0.15 - 0.10 \times (1.48 - 1.40) = 0.15 - .08 = 0.142,$$

which is 0.02 above the initial rate (instead of 0.10, if unhedged).

In the example, setting beta equal to 0.10 clearly lowers the risk. The standard approach is to choose β so as to make the variance of the hedged cash flow as small as possible. But we already know the solution. If we had written the problem as one of minimizing $\text{var}(\tilde{\epsilon})$ where $\tilde{\epsilon} := \tilde{y} - \beta\tilde{x}$, you would immediately have recognized this to be a "regression" problem, with the usual regression beta as the solution:

$$\begin{aligned} \beta &= \text{the slope coefficient from } \tilde{S}^{(e)}_{T_1} = \alpha + \beta \tilde{f}^{(h)}_{T_1,T_2} + \tilde{\epsilon}, \\ &= \frac{\text{cov}(\tilde{f}^{(h)}_{T_1,T_2}, \tilde{S}^{(e)}_{T_1})}{\text{var}(\tilde{f}^{(h)}_{T_1,T_2})}. \end{aligned} \tag{6.3}$$

DoItYourself problem 6.2

Formally derive this result. First write out the variance of the hedged cash flow for a given β, using the fact that the (known) current futures price does not add to the variance. Then find the value for β that minimizes the variance of the remaining risk.

We now look at a number of special cases.

6.4.2 Case 1: The Perfect Match

There is a perfect match if the futures contract expires at T_1 (that is, $T_2 = T_1$) and $e = h$. For example, assume there is a SEK contract with exactly the same date as your exposure. The convergence property means that $\tilde{f}^{(e)}_{T_1,T_2} = S^{(e)}_{T_1}$: on the last day of trading a SEK futures price exactly equals the spot rate at the same moment because both stipulate delivery at $t+2$. Thus, in this special case of a perfect match, Equation [6.3] tells us, we should regress the variable upon itself. There of course is no need to actually do so: in that regression, the slope coefficient (and the R^2) can only be unity. So you sell forward one for one: if the exposure is B units of forex, you sell B units. In short, this is standard hedging where nothing needs to be estimated.

But usually one is not that lucky:

6.4.3 Case 2: The Currency-Mismatch Hedge or Cross-Hedge

We now consider a case where the futures contract matches the maturity of the foreign-currency inflow but not the currency (h $\neq$ e). For instance, the US exporter's SEK inflow is hedged using a EUR future. We can use the convergence property $f_{T_1,T_1} = S_{T_1}$ to specify the hedge ratio as

$$\begin{aligned} \beta &= \frac{\text{cov}(\tilde{S}^{(e)}_{T_1}, \tilde{S}^{(h)}_{T_1})}{\text{var}(\tilde{S}^{(h)}_{T_1})}, && (6.4) \\ &= \text{the slope coefficient in } \tilde{S}^{(e)}_{T_1} = \alpha + \beta\, \tilde{S}^{(h)}_{T_1} + \tilde{\epsilon}. && (6.5) \end{aligned}$$

This measure of linear exposure will come up again and again in this book, most prominently in Chapter 9 on option pricing and hedging, or in Chapter 13 when we quantify operating exposure, so reading on is useful even if we get technical, initially. Actually, another reason for getting technical is that it helps us understand the pros and cons of the relative-value hedging rule that we introduced before, like hedging every SEK by 0.10 EUR, the current cross-rate.

Recall that in the definition of exposure we hold the time constant, and we instead compare possible future scenarios. Similarly, our regression is, in principle, forward looking: it should be run across a representative number of (probability-weighted) possible future scenarios. This is not easy, so you may want to run the regression on past data instead. One assumption then is that β is constant, so that the past is a good guide to the future. For technical and statistical reasons that are beyond the scope of this chapter, one should not regress levels of exchange rates on levels if the data are time series. A regression between changes of the variables, in contrast, would be statistically more acceptable:

$$\text{regress } \Delta S^{(e)}_t = \alpha' + \beta\, \Delta S^{(h)}_t + \tilde{\epsilon}'_t, \qquad (6.6)$$

where, this time, deltas refer to changes over time. Many careful researchers would still be unhappy with this, and actually prefer to work with a regression in percentage changes: in a long time series with much variation in the level of S, it is hard to believe that de distribution of ΔS is constant. So if we use s as shorthand for $\Delta S/S$, we would use the following equation:

$$\tilde{s}_t^{(e)} = \alpha'' + \gamma\, \tilde{s}_t^{(h)} + \tilde{\epsilon}_t''. \tag{6.7}$$

The assumption now is that γ is constant, not β. If you run a regression between percentages, you need to transform the slope γ from an elasticity into a partial derivative:[6]

$$\begin{aligned}\text{if } \tilde{s}_t^{(e)} = \alpha'' + \gamma\, \tilde{s}_t^{(h)} + \tilde{\epsilon}_t'' \text{ then } \beta &= \gamma\, \frac{S_t^{(e)}}{S_t^{(h)}} &(6.8)\\ &= \gamma\, S_t^{(cross)}, &(6.9)\end{aligned}$$

where $S^{(cross)}$ is the cross rate. In our example,

$$\frac{S_t^{(e)}}{S_t^{(h)}} \text{ has dimension } \frac{\text{USD/SEK}}{\text{USD/EUR}} = \text{EUR/SEK}, \tag{6.10}$$

so the cross rate is the value of one SEK in EUR, which is euros per crown or, generally, h/e.

If γ is unity, we get the relative-value rule that we started out with in the first example, where we hedged each crown with 0.10 euros because the initial cross rate is EUR/SEK 0.10. This is a rule that practitioners often use. They do not actually run this regression: instead, they just guess that the gamma equals unity. For instance, Boston SC's expects that every percentage in the EUR (against the USD) on average leads to a similar change in the USD value of the SEK. [This is slightly more general than our earlier story of a fixed cross rate: now each percentage change in the euro's value is assumed to lead to the same percentage change in the crown's value *on average*; sometimes it may be more, sometimes less, but on average the change is the same.] Then the hedge ratio would simply be set equal to the cross rate:

$$\text{Rule of thumb for cross hedge: } \gamma = 1 \text{ so } \beta = S_t^{(h/e)}. \tag{6.11}$$

Example 6.12

Again assume spot rates of 1.40 for the EUR and 0.14 for the SEK. The quick-and-dirty hedge ratio would be set equal to the cross rate, the value of on SEK in EUR, which equals $0.14/1.40 = 0.10$. The reason is that you think that percentage changes of the two currencies will be similar ($\gamma = 1$), but since the EUR is worth about ten Kronar now, one EUR would change by as much as would ten Kronar. Therefore,

[6] In terms of a regression of y on x, the exposure is written as $\frac{\Delta y}{\Delta x}$. An elasticity equals $e = \frac{\Delta y}{\Delta x} \times \frac{x}{y}$, so $\frac{\Delta y}{\Delta x} = e \times \frac{y}{x}$.

one Euro shorted would hedge about ten SEK. In other words, 0.10 Euros per SEK will do.

Suppose, alternatively, that you prefer to run a regression between monthly percentage changes on SEK and EUR, and the slope is 0.96 with an R^2 of 0.864. Then

$$\text{regression-based hedge ratio} = 0.96 \times 0.10 = 0.096.$$

That is, you'd lower your hedge ratio.

The rule of thumb is almost surely biased, which is bad, but has one big advantage: it has zero sampling error. Let us explain each statement. First, the assumption of unit gamma's across the board does not make sense, statistically. For example, if it *were* true, then the reverse regression, between EUR and SEK rather than the inverse, would also produce a unit gamma, but this is mathematically possible only when there is no noise. It is easy to verify that the product of the two gamma's—the one from y on x and the one from x on y— is the R^2, which is surely a number below unity; so one expects at least one of the two gamma's to be below unity, and normally both will be below unity. I elaborate on this in Teknote 6.1.

But while the drawback of the rule of thumb is a bias, it has the advantage of no sampling error. If you actually run regressions, then the estimated sample will randomly deviate, depending on sampling coincidences, even if nothing structural has changed. Now from the point of view of the user, sampling error is as bad as bias. For instance, if the true gamma (known to the Great Statistician in the Sky only) is 0.95, then the error introduced by an estimated gamma of 0.90 is as bad as the bias introduced by the rule-of-thumb value, unity. Likewise, hedging with a unit gamma would be as bad as hedging with an estimated gamma that equals, with equal probability, 1.00 or 0.90. So it all depends on squared bias versus estimation variance. Experiments (Sercu and Wu, 2000) show that the rule of thumb does better than the regression-based hedge if the relation between i and j is close, which is the case for the USD/SEK and USD/EUR rates. When the link between the two variables becomes lower, sampling-error variance increases but so does the bias, and in fact bias tends to become the worse of the two evils.

6.4.4 Case 3: The Delta hedge

Suppose now that there is a SEK contract, but for the wrong date instead of for the wrong currency. Our money comes in Feb 15, while the contract expires March 20, for example. So our futures contract will still have a 35-day remaining life when it is liquidated. In principle we'd have to regress possible spot values for the SEK on the corresponding 35-day futures price of the SEK. One problem is that we do not have time-series data on 35-day futures: the real-world data have a daily-changing maturity.

There are two ways out, both connected to IRP. Since futures are almost indistinguishible from forwards, we know that

$$\tilde{f}^{(e)}_{T_1,T_2} \approx \tilde{S}^{(e)}_{T_1} \frac{1+\tilde{r}_{T_1,T_2}}{1+\tilde{r}^{(e)}_{T_1,T_2}}, \tag{6.12}$$

where the risk-free rates r now get tildes because we do not yet know what they will be, on Feb 15. So one way to solve the ever-changing-maturity problem in the data is to construct forward rates from spot and interest data, probably using 30-day *p.a.* rates to approximate the 35 *p.a.* data.[7] The other way out is to use a rule of thumb. Inverting Equation [6.12], we get

$$\tilde{S}^{(e)}_{T_1} = \frac{1+\tilde{r}^{(e)}_{T_1,T_2}}{1+\tilde{r}_{T_1,T_2}} \tilde{f}^{(e)}_{T_1,T_2}. \tag{6.13}$$

The rule of thumb then follows under a not very harmful assumption, namely that there is no uncertainty about the interest rates. For instance, suppose you knew that the ratio $(1+r^{(e)})/(1+r)$ 35 days would be 1.005 on Feb 15. Then Equation [6.12] would specialize into

$$\tilde{S}^{(e)}_{T_1} = 1.005\tilde{f}^{(e)}_{T_1,T_2}, \tag{6.14}$$

which tells us immediately that the forward-looking regression coefficient of $S^{(e)}$ on $f^{(e)}$ is 1.005. So the rule of thumb for the delta hedge is to set the hedge ratio equal to the forecasted ratio $(1+r^{(e)})/(1+r)$ 35 days for Feb 15. Experiments show that it hardly matters how you implement this: take the current 35-day rates, or forecasts implicit on forward interest rates (if available). Also, since the regression (if you would run it) has a very high R^2, the bias is tiny and the rule of thumb does quite well.

6.4.5 Case 4: The Cross-and-Delta hedge

Now combine the problems: we use a EUR contract expiring March 20 to hedge SEK that come forth on Feb 15. In principle we have to regress possible SEK spot rates on 35-day EUR futures.

The rule of thumb is a combination of the two preceding ones: set the hedge ratio equal to the current cross rate times the forecast ratio $(1+r^{(e)})/(1+r)$. Again, the rule of thumb does quite well when the currencies i and j are closely related and the R^2, therefore, is high.

[7] This would also solve synchronization problems in data: spot and interest rates are observed at the same time, while the futures prices may be from a different data base and observed at a different time of the day.

6.4.6 Adjusting for the Sizes of the Spot Exposure and the Futures Contract

Thus far, we have assumed that the exposure was one unit of a first foreign currency, e, and that the size of one futures contract is one unit of another foreign currency, h. If the exposure is a larger number, say n_e, then the number of contracts one needs to sell obviously goes up proportionally, while if the size of the futures contract is n_h rather than unity, the number of futures contracts goes down proportionally. Thus, the generalized result is as follows: the number of contracts to be sold in order to hedge n_e units of currency j using a futures contract with size n_h units of currency i is given by

$$\text{hedge ratio } = \frac{n_e}{n_h}\beta, \tag{6.15}$$

where β can be regression-based or a rule-of-thumb number.

Example 6.13

Suppose that you consider hedging a SEK 2.17m inflow using EUR futures with a contract size of EUR 125,000. A regression based on 52 points of weekly data produces the following output:

$$\Delta S_{[\text{USD/SEK}]} = 0.003 + 0.105 \Delta f_{[\text{USD/EUR}]}. \tag{6.16}$$

with an R^2 of 0.83 and a t-statistic of 15.62. Then:

- In light of the high t-statistic, we are sure that there actually is a correlation between the USD/SEK spot rate and the USD/EUR futures price.
- Assuming all correlation between the two currencies is purely contemporaneous, hedging reduces the total uncertainty about the position being hedged by an estimated 83 percent. If the horizon is more than one week and if there are lead-lag reactions between the currencies, this estimate is probably too pessimistic.
- The regression-based estimated for the number of contracts to be sold is

$$\text{hedge ratio } = \frac{2,170,000}{125,000} \times 0.105 = 1.822, \tag{6.17}$$

or, after rounding, 2 contracts.

6.4.7 More About Regression-based Hedges

When implementing a regression-based hedge you need to think about a number of items:

- **Estimation error** Novices think of a regression coefficient as a sophisticated number computed by clever people. Old hands dejectedly look at the huge error margin, conveniently calculated for you by the computer program, and then sink into even deeper sloughs of despond despair when they remember that the calculated margin is almost surely too optimistic: the real world is never so simple as our computers assume.

- **Errors in the regressor** If you use futures data, there is a problem of bid-ask noise (you would probably like to have the midpoint rate, but the last traded price is either a bid or an ask—you don't know which), changing maturities, jumps in the basis when the data from an expiring short contract are followed by prices from a 3-month one, and synchronization problems between spot and forward prices. So if you use futures transaction data there is an errors-in-variables problem that biases the β estimate towards zero.

  Many of these problems can be solved by using forward prices computed from midpoint spot and money-market rates for the desired maturity $T_2 - T_1$.

- **Lead/lag reactions and the intervalling effect** The SEK tends to stay close to the EUR, from an USD perspective. But this means that, if the EUR appreciates, for example, and the SEK does not entirely follow during the same period, then there typically is some catching-up going on in the next period. This means that the correlation between changes in the Euro and the Krona is not purely contemporaneous.

  This gives rise to the *intervalling effect.* The beta computed from, say, five-minute changes is quite low, but the estimates tend to increase if one goes to hourly, daily, weekly, and monthly intervals. This is because, the longer the interval, the more of the lagged reaction is captured within the interval.

  Example 6.14

  Suppose that, "in the long run" every percentage change in the EUR means an equal change in the SEK's value, but only three-quarter of that takes place the same day, with the rest taking place the next day, on average. Then your estimated γ from daily data would be more like 0.75 than 1.00 as your computer overlooks the non-contemporaneous linkages. But if you work with weekly data (five trading days), then for four of the days the lagged effect is included into the same week and picked up by the covariance; only 0.25 of the last-day effect is missed, out of 5 days' effects, causing a bias of just $0.25/5 = 0.05$. Obviously, with monthly data the problem is even smaller.

  The intervalling effect means that, ideally, the interval in your regression should be equal to your hedging horizon, otherwise the beta tends to be way too low. This can be implemented in three ways. First, you could take *non-overlapping holding periods.* The problem is that this often leaves you with too few useful observations. For instance, if your horizon $T_2 - T_1$ equals three months and you think that data older than 5 year are no longer relevant, you have just a pitiful 20 quarterly observations. Second, you could use *overlapping observation periods.* For example, you work with 13-week periods, the first covering weeks 1-13, the next weeks 2-14, etc. This leaves you more useful information; but remember that the usual R^2 and t-statistics are no longer reliable because of the overlap created between the observations. (Hansen and Hodrick (1980) show you how to adjust the confidence intervals.) Third, you could use a clever, *non-standard regression* technique that tries to capture the relevant lead/lag affects. Examples are the instrumental-variables estimators by Scholes and Williams (1977) or Sercu, Vandebroek and Vinaimont (2008), or the multivariate-based beta by Dimson (1979), or an error-correction model like Sultan and Kroner (1993).

6.4.8 Hedging with Futures Using Contracts on More than One Currency

Occasionally one uses more than one futures contract to hedge. For instance, a US hedger exposed to NOK may want to use EUR and GBP contracts to get as close as possible to the missing NOK contract. In principle, the solution is to regress NOK spot prices on EUR *and* GBP futures prices, and use the multiple regression coefficients as hedge ratios. Rules of thumb do not exist here. If one uses actual regression of past time-series data, one would of course resort to first changes (ΔS and Δf.) or percentage changes.

This finishes our discussion of how to adjust the size of the hedge position for maturity and currency mismatches. In the appendix we digress on interest-rate futures—not strictly an international-finance contract, but one that is close to the FRA's discussed in an appendix to Chapter 4 which, you will remember, are very related to currency forwards and forward-forward swaps. We conclude with a discussion of how forwards and futures can co-exist. Clearly each must have its own important strengths, otherwise one of them would have driven out the other.

6.5 The CFO's conclusion: Pros and Cons of Futures Contracts Relative to Forward Contracts

Now that we understand the differences between futures and forwards, let us compare the advantages and disadvantages of using futures rather than forwards. The **advantages** of using futures include:

- Because of the institutional arrangements in futures markets, the default risk of futures contracts is low. As a consequence, relatively unknown players without an established reputation or without the ability to put up substantial margin can trade in futures markets. This is especially relevant for speculators who are not interested in actual delivery at maturity.

- Because of standardization, futures markets have low transaction costs; commissions in futures markets tend to be lower than in forward markets, especially for small lot sizes. Remember that to get wholesale conditions in the forward market, one needs to deal in millions of USD, while in the futures section 100,000 or thereabouts suffices.

- Given the liquidity of the secondary market for futures, futures positions can be closed out early with greater ease than forward contracts.

Clearly, there are also **drawbacks** to futures contracts—otherwise, forward markets would have disappeared entirely:

- One drawback is the standardization of the futures contract. A creditworthy hedger has to choose between an imperfect but cheap hedge in the futures markets and a more expensive but exact hedge in the forward market. The standardization of the futures contracts means that one will rarely be able to find a contract of exactly the right size or the exact same maturity as that of the underlying position to be hedged.

- Futures contracts exist only for a few high-turnover exchange rates. This is because futures markets cannot survive without large trading volumes. Thus, for most exchange rates, a hedger has to choose between forward contracts or money-market hedges, or a cross-hedge in the futures markets. A cross-hedge is less effective because the relationship between the currency one is exposed to and the currency used as a hedge instrument is obscured by cross-rate risk.

- Also, marking to market may create ruin risk for a hedger. A firm that expects to receive EUR 100m nine months from now faces no inflows or outflows when it hedges in the forward market. In contrast, the daily marking to market of a futures contract can create severe short-term cash flow problems. It is not obvious that interim cash outflows can always be financed easily.

- Assuming that financing of the interim cash flows is easy, marking to market still creates interest-rate risk. The daily cash flows must be financed/deposited in the money markets at interest rates that are not known when the hedge is set up. This risk is not present in forward hedging. The correlation between futures prices and interest rates is typically rather low, implying that the interest rate risk is small *on average*; but in an individual investment the *ex post* effect can be larger.

- Lastly, futures markets are available only for short maturities. Maturities rarely exceed eleven months, and the markets are often thin for maturities exceeding six months. In contrast, forward contracts are readily available for maturities of up to one year, and today the quotes for forward contracts extend up to ten years and more.

We see that the competing instruments, forwards and futures, appear to cater to two different clienteles. As a general rule, forward markets are used primarily by corporate hedgers, while futures markets tend to be preferred by speculators. Mark these words: it's a general rule, not an exact law.

6.6 Appendix: Eurocurrency Futures Contracts

Eurocurrency futures contracts can be used to hedge or to speculate on interest risk, in contrast to currency futures, which allow one to hedge (or speculate on) exchange risk. That is, eurocurrency futures are the futures-style counterparts of Forward Forward contracts and Forward Rate Agreements, in the same way as futures contracts on currencies relate to currency forward contracts.

The first traded eurocurrency futures contract was the eurodollar contract traded at the International Money Market on the Chicago Mercantile Exchange (CME), now working on a merger with its arch-rival commodity exchange, the Chicago Board of Trade (CBOT). Eurodollar futures were quickly introduced also in the London International Financial Futures Exchange (LIFFE), now part of Euronext, and the Singapore Monetary Exchange (SIMEX). Currently, most financial centers of countries with a well-developed capital market have a contract written on the local interbank interest rate—for instance, the EUR contract that used to be traded on the *Marché à Terme International de France* (MATIF) in Paris, now part of Euronext's LIFFE CONNECT. As can be seen from Table 6.2, many exchanges also trade a few foreign contracts—for instance, JPY in SIMEX. (The list is just a sample; no completeness is intended.)

Figure 6.2: **Some Interest-futures Markets**

	Underlying	Exchange	Contract size*	longest**
AUD	90-day accepted bills	SFX	500,000	3y
BEF	3-month BIBOR	BELFOX	25,000,000	9m
CAD	Canadian B/A	ME	1,000,000	2y
DEM	3-month LIBOR	LIFFE, MATIF, DTB	1,000,000	9m
EIP	3-month DIBOR	IFOX	100,000	9m
BRC	Domestic CD	BM&F	10,000	11m
GBP	3-month Euro-sterling	LIFFE	500,000	9m
JPY	3-month LIBOR	TIFFE, SIMEX	1,000,000	9m
FRF	3-month PIBOR	MATIF	5,000,000	9m
NZD	90-day accepted bills	NZFE	500,000	2y
USD	3-month LIBOR	CME, SIMEX, LIFFE,TIFFE	1,000,000	2y
USD	1-month LIBOR	CME, CBOT	3,000,000	2y
USD	30-day Federal Funds	CBOT	5,000,000	2y

*: at first exchange listed. Contract size at other exchanges may differ
**: life of longest contract, at first exchange listed; m=month, y=year

SFX: Sidney Futures Exchange; BELFOX: Belgian Futures and Oprions Exchange; ME: Montreal Exchange; LIFFE: London International Financial Futures Exchange; MATIF: Marché à Terme International de France; DTB: Deutsche TerminBörse; IFOX: Irish Futures and Options Exchange; BM&F: Bolsa y Mercantil y de Futuros (Saõ Paulo); TIFFE: Tokyo International Financial Futures Exchange; SIMEX: Singapore Monetary Exchange; NZFE: New Zealand Futures Exchange; CME: Chicago Mercantile Exchange (includes IMM); CBOT: Chicago Board of Trade.

Source Data from Sercu and Uppal, R., International Financial Markets and the Firm, South-Western Publishing Company, 1995.

Many of the European futures contracts are in effect *collateralized forward contracts*, where the investor puts up more collateral (securities, or interest-bearing deposits) if the price evolution is unfavorable, rather than making a true payment. As was explained in Chapter 6, a collateralized forward contract is not subject to interest risk.

Let us now see how a eurocurrency futures contract works. A useful first analogy is to think of such a contract as similar to a futures contract on a CD, where the expiration day, T_1, of the futures contract precedes the maturity date, T_2, of the CD by, typically, three months. (The three-month money-market rate is widely viewed as the representative short-term rate,) Thus, such a futures contract serves to lock in a three-month interest rate at time T_1.

Example 6.15

Suppose that in January you agree to buy, in mid-March, a CD that expires in mid-June. The maturity value of the CD is 100, and the price you agree to pay is 99. This means that the return you will realize on the CD during the last three months of its life is $(100 - 99)/99 = 1.0101$ percent, or 4.0404 percent simple interest on a yearly basis. Thus, this forward contract is analogous to signing an FRA at 4.0404 percent *p.a.* for three months, starting mid-March.

In the example, we described the futures contract as if it were a forward contract. If there is marking to market, the interest risk stemming from the uncertain marking-to-market cash flows will affect the pricing. Another complication with futures is that the quoted price is often different from the effective price, as we discuss below. Still, it helps to have the above example in mind to keep from getting lost in the institutional details. We first derive how forward prices on T-bills or CDs are set, and how they are linked to the forward interest rate. We then discuss the practical problems with such a system of quotation, and explain how this has led to a modern futures quote, an animal that differs substantially from the forward price on a T-bill or CD.

6.6.1 The Forward Price on a CD

The forward price on a CD is just the face value (1, most often quoted as 100 percent) discounted at the forward rate of return, r^f_{t,T_1,T_2}. To understand this property, consider a forward contract that expires at T_1 and whose underlying asset is a euro-CD maturing at T_2 ($> T_1$). Since the euro-CD has no coupons, its current spot price is:

$$V_t = \frac{1}{1 + r_{t,T_2}}, \tag{6.18}$$

where, as always in this textbook, r_{t,T_2} denotes an effective return, not a *p.a.* interest rate. The CD's forward price at t, for delivery at T_1, is this spot value grossed up with the effective interest between t and T_1 (line 1, below), and the combination of the two spot rates then gives us the link with the forward rate:

$$\begin{aligned} V^f_{t,T_1,T_2} &= V_t\,(1+r_{t,T_1}) \\ &= \frac{1+r_{t,T_1}}{1+r_{t,T_2}}, && \text{see [6.18]}, \\ &= \frac{1}{1+r^f_{t,T_1,T_2}}, && \text{see [4.48]}. \end{aligned} \tag{6.19}$$

Example 6.16

Consider a six-month forward on a nine-month bill with face value USD 1. Let the *p.a.* interest rates be 4 percent for nine months, and 3.9 percent for six months. Then $r_{t,T_2} = (9/12) \times 4\% = 0.03$, so that the spot price (quoted as a percentage) is equal to:

$$V_t = \frac{100\%}{1.03} = 97.087\%. \tag{6.20}$$

Also, $r_{t,T_1} = (6/12) \times 3.9\% = 0.0195$; thus, the forward price today is:

$$V^f_{t,T_1,T_2} = 97.087 \times 1.0195 = 98.981\%. \tag{6.21}$$

Alternatively, we can compute the six-month forward price on a nine-month T-bill via the forward rate of return:

$$1 + r^f_{t,T_1,T_2} = \frac{1.03}{1.0195} = 1.010,299 \Rightarrow V^f_{t,T_1,T_2} = \frac{1}{1.010,299} = 98,981\%. \tag{6.22}$$

For some time, interest rate futures markets in Sydney were based on this system of forward prices for CDs. Although the system is perfectly logical, traders and investors are not fond of quoting prices in this way. One reason is that traders and dealers are more familiar with *p.a.* interest rates than with forward prices for deposits or CDs. The process of translating the forward interest rate into a forward price is somewhat laborious: Equation [6.19] tells us that the unfortunate trader has to divide the per annum forward rate by four, add unity, and take the inverse to compute the normal forward price as the basis for trading. A second problem is that real-world interest rates are typically rounded to one basis point (0.01 percent). Thus, unless forward prices are also rounded, marking to market will result in odd amounts. These very practical considerations lead to a more user-friendly manner of quoting prices for futures on CDs.

6.6.2 Modern Eurodollar Futures Quotes

To make life easier for the traders, rather than quoting a true futures price, most exchanges quote three-month eurodollar futures contract prices as:

$$\text{Quote} = 100 - [\text{per annum forward interest rate}], \tag{6.23}$$

Figure 6.3: **Interest-rate Futures as reported by the Financial Times**

06/06/06		Open	Sett	Change	High	Low	Est.Vol	Open int
Euribor 3m*	Jun	96.97	96.96	-0.01	96.98	96.96	98.491	553.392
Euribor 3m*	Sep	96.69	96.69	-1.01	96.70	96.68	107.380	582.722
Euribor 3m*	Dec	96.49	96.48	-0.02	96.50	96.47	110.900	608.862
Euribor 3m*	Mar	96.39	96.37	-0.03	96.40	96.36	95.601	487.298
Euribor 3m*	Jun	96.30	96.29	-0.03	96.32	96.28	86.201	414.737
...								...
Euroswiss 3m*	Jun	98.48	98.47	-0.01	98.48	98.46	8.045	75.781
...								...
Sterling 3m*	Jun	95.27	95.26	-0.02	95.27	95.26	15,291	416,929
...								...
Eurodollar 3m#	Jun	94.68	94.67	-0.01	94.69	94.66	117,026	1297,150
...								...
FedFnds 30d+	Jun	94.970	94.965	-0.005	94.970	94.965	1,078	107,907
...								...
Euroyen 3m±	Jun	99.655	99.650	-0.020	99.665	99.650	15.364	586.998
...								...

Contracts are based on volumes traded in 2004 Sources: * LIFFE # CME + CBOT ± TIFFE

and base the marking to market on *one-fourth* of the change in the quote.

Before we explain the marking-to-market rule, let us first consider the quotation rule given in Equation [6.23]. This quote decreases when the forward interest rate increases—just as a true forward price on a T-bill—and the long side of the contract is still defined as the one that wins when the quote goes up, the normal convention in futures or forward markets. However, one major advantage of this price-quoting convention is that a trader or investor can make instant decisions on the basis of available forward interest quotes, without any additional computations.

Example 6.17

Let the *p.a.* forward interest rate be 4.1 percent *p.a.* for a three-month deposit starting at T_1. The true forward price would have been computed as

$$V^f_{t,T_1,T_2} = \frac{1}{1 + (1/4)\,0.041} = 98.985,300 \approx 98,99. \tag{6.24}$$

In contrast, the eurodollar forward quote can be found immediately as 100 percent – 4.1 percent = 95,9 percent.

The second advantage of the "100 minus interest" way of quoting is that such quotes are, automatically, multiples of one basis point because interest rates are multiples of one basis point. With a standard contract size of USD 1m, one tick (equal to 1/100 of a percent) in the interest rate leads to a tick of 1m × 0.0001 = USD 100 dollars in the underlying quote (no odd amounts here). Note that, since marking to market is based on one-fourth of the change in the quote, a one-tick change in the interest rate leads to a USD 25 change in the required margin.

To understand why marking to market is based on one-fourth of the change of the quote, go back to the correct forward price, Equation [6.19]. The idea is to undo the fact that the change in the quote (Equation [6.23]) is about four times the change in the correct forward price (Equation [6.19]). To understand this, note that $T_2 - T_1$

corresponds to three months (1/4 year). Thus, as a first-order approximation,

$$\frac{1}{1+r^f_{..}} \approx 1 - r^f_{..} = 1 - 1/4 \times (4r^f_{..}) = 1 - 1/4 \times [\text{p.a. forward interest rate}] \quad (6.25)$$

Thus, the change in the true forward price is about one-fourth of the change in the futures quote. To bring marking to market more or less in line with normal (price-based) contracts, the changes in the quote (or in the *p.a.* forward interest rate) must be divided by four. If this were not done, a USD 1m contract would, in fact, hedge a deposit of roughly USD 4m, which would have been very confusing for novice buyers and sellers.

Example 6.18

Suppose that you hold a five-month, USD 1m CD and you want to hedge this position against interest rate risk two months from now. If, two months from now, the three-month interest rate drops from 4 percent to 3.9 percent, the market value of your deposit increases from 1m/(1 + (1/4) 0.04) = 990,099.01 to 1m/(1 + (1/4) 0.039) = 990,344.14, a gain of USD 245.13. The price quoted for a futures contract would change by 0.1 percent or, on a USD 1m contract, by USD 1,000. Marking to market, however, is one-fourth of that, or USD 250. So the marking-to-market cash flows on the eurodollar futures contract would reasonably match the 245.13-dollar change in the deposit's market value.

The pros and cons of interest futures, as compared to FRAs, are the same as for any other futures contract. The main advantage is an active secondary market where the contract can be liquidated at any moment, and the lowish entry barriers because of the efficient way of handling security: ask for cash only if and when it is needed. But that cannot be the end of the story. Also FRAs have some advantages over T-bill futures and bond futures, and these advantages are similar to those of forward exchange contracts over currency futures contracts, as discussed in Chapter 6.

- FFs or FRAs are pure forward contracts, which means that there is no marking to market. It follows that, by using FFs or FRAs, one avoids the additional interest risk that arises from marking to market.

- In the absence of marking to market, there is no ruin risk. The firm need not worry about potential cash outflows that may lead to liquidity problems and insolvency.

- In the absence of marking to market, there is an exact arbitrage relationship between spot rates and forward rates; hence these contracts are easy to value. In contrast, the pricing of a futures-style contract is more difficult because of interest risk—covariance between market values and the interest-rate evolution, which in the case of interest derivatives is, of course, stronger than for futures on currency or stocks of commodities.

- FRAs are tailor-made, over-the-counter instruments and are, therefore, more flexible than (standardized) futures contracts. Hedgers with small exposures may not like a contract of USD 1m, and if three-month futures are used to hedge against a change in the four-month or nine-month interest rate, the hedge is, at best, imperfect.

- The menu of underlyings is quite limited: three-month rates, and (in the bond market, which we have not discussed) medium-term bonds.

For these reasons, FFs and FRAs are better suited for arbitrage or hedging than are futures.

Technical Note 6.1 Why gammas are below unity

Technically, the rule of thumb of setting γ equal to unity is supposedly based on the assumption that no change in the relative values is expected. So if the percentage changes in the SEK and EUR spot rates are denoted by y and x, respectively, then the trader's feeling is that $\mathrm{E}_t(\tilde{y}) = \mathrm{E}_t(\tilde{x})$. But this is an unconditional statement. A regression is a conditional statement: what do we expect about $\tilde{y}$ for a given value of x. Suppose, for instance, that both $\tilde{x}$ and $\tilde{y}$ have an unconditional mean of zero. Then in the regression $y = a + bx + e$ the slope b can be unity indeed—but it can also be 0.5, or zero, or –1 for the matter. Indeed, if $y = a + bx + e$ holds, then $\mathrm{E}(\tilde{y}) = a + b\mathrm{E}(\tilde{x}) + 0$ follows, and since the expectations are zero, the only constraint is that a must be zero; the slope b can still be anything.

Thus, conditional and unconditional expectations are different animals. In our case,

$$\mathrm{E}_t(\tilde{s}_T^{sek}) = \mathrm{E}_t(\tilde{s}_T^{eur}) \text{ does not imply } \mathrm{E}_t(\tilde{s}_T^{sek}|s_T^{eur}) = s_T^{eur}, \tag{6.26}$$

even though the reverse statement does hold:

$$\mathrm{E}_t(\tilde{s}_T^{sek}|s_T^{eur}) = s_T^{eur} \text{ does imply } \mathrm{E}_t(\tilde{s}_T^{sek}) = \mathrm{E}_t(\tilde{s}_T^{eur}). \tag{6.27}$$

The above uses statistics to make the point, which may fail to impress many readers; so let's also think of the economics. Exchange rates in our currency threesome can move because there is news about the US, or about Euroland, or about Sweden. A lot of world news has implications for all three, but some news is purely local—for instance, housing starts in Sweden may be below what pundits had expected while things are fine elsewhere.

Suppose the USD/EUR rate increases. This could be because of relatively bad news about the US, or relatively good news about Euroland. If the source is pure dollar news, then also the USD/SEK rate would go up by a similar percentage, as there is no reason for the EUR/SEK rate to change: it's the dollar that falls, not the euro that rises. But if, in contrast, the source is pure euro news, then the appreciation of the USD/EUR rate is because the euro rises not because the dollar falls, meaning that the USD/SEK rate would not budge. To sum up, in our stylized story,

a. if there's dollar news, then $\mathrm{E}_t(\tilde{s}_T^{sek}|s_T^{eur}) = s_T^{eur}$ (the crown rises as much as the euro)

b. if there's euro news, then $\mathrm{E}_t(\tilde{s}_T^{sek}|s_T^{eur}) = 0$ (the crown does not follow the euro)

and since we don't know whether the news will be about the US or about Euroland, gamma must be between unity (case a.) and zero (case b.). Where exactly gamma would be depends on the relative probabilities of either type of news, and also about how earth-moving the news tends to be. For instance, if both types of news are equally likely but european news merely rises eyebrows while US news causes heart attacks, the first scenario would dominate the distribution and gamma would be closer to unity than to zero.

6.7 Test Your Understanding

6.7.1 Quiz Questions

1. For each pair shown below, which of the two describes a forward contract? Which describes a futures contract?

   (a) standardized/made to order
   (b) interest rate risk/no interest rate risk
   (c) ruin risk/no ruin risk
   (d) short maturities/even shorter maturities
   (e) no secondary market/liquid secondary market
   (f) for hedgers/speculators
   (g) more expensive/less expensive
   (h) no credit risk/credit risk
   (i) organized market/no organized market

2. Match the vocabulary below with the following statements.

(1) organized market	(11) maintenance margin
(2) standardized contract	(12) margin call
(3) standardized expiration	(13) variation margin
(4) clearing corporation	(14) open interest
(5) daily recontracting	(15) interest rate risk
(6) marking to market	(16) cross-hedge
(7) convergence	(17) delta-hedge
(8) settlement price	(18) delta-cross-hedge
(9) default risk of a future	(19) ruin risk
(10) initial margin	

   (a) Daily payment of the change in a forward or futures price.
   (b) The collateral deposited as a guarantee when a futures position is opened.
   (c) Daily payment of the discounted change in a forward price.
   (d) The minimum level of collateral on deposit as a guarantee for a futures position.
   (e) A hedge on a currency for which no futures contracts exist and for an expiration other than what the buyer or seller of the contract desires.
   (f) An additional deposit of collateral for a margin account that has fallen below its maintenance level.
   (g) A contract for a standardized number of units of a good to be delivered at a standardized date.

(h) A hedge on foreign currency accounts receivable or accounts payable that is due on a day other than the third Wednesday of March, June, September, or December.

(i) The number of outstanding contracts for a given type of futures.

(j) The one-day futures price change.

(k) A proxy for the closing price that is used to ensure that a futures price is not manipulated.

(l) Generally, the last Wednesday of March, June, September, or December.

(m) Organization that acts as a "go-between" for buyers and sellers of futures contracts.

(n) The risk that the interim cash flows must be invested or borrowed at an unfavorable interest rate.

(o) A hedge on a currency for which no futures contract exists.

(p) The risk that the price of a futures contract drops (rises) so far that the purchaser (seller) has severe short-term cash flow problems due to marking to market.

(q) The property whereby the futures equals the spot price at expiration.

(r) Centralized market (either an exchange or a computer system) where supply and demand are matched.

The table below is an excerpt of futures prices from an old *The Wall Street Journal* copy. Use this table to answer Questions 3 through 6.

						Lifetime		Open
	Open	High	Low	Settle	Change	High	Low	Interest
JAPAN YEN (CME) — 12.5 million yen; $ per yen (.00)								
June	.9432	.9460	.9427	.9459	+.0007	.9945	.8540	48.189
Sept	.9482	.9513	.9482	.9510	+.0007	.9900	.8942	1,782
Dec	.9550	.9610	.9547	.9566	+.0008	.9810	.9525	384
	Est vol 13,640; vol Fri 15,017; open int 50,355, +414							
New Zealand Dollar (CME) — 125,000 dollars; $ per dollar								
June	.5855	.5893	.5847	.5888	+.0018	.6162	.5607	87,662
Sept	.5840	.5874	.5830	.5871	+.0018	.6130	.5600	2,645
Dec	.5830	.5860	.5830	.5864	+.0018	.5910	.5590	114
	Est vol 40,488; vol Fri 43,717; open int 90,412, -1,231							
Swiss Franc (CME) — 100,000 francs; $ per franc								
June	.7296	.7329	.7296	.7313	+.0021	.7805	.7290	43,132
Sept	.7293	.7310	.7290	.7297	+.0018	.7740	.7276	962
Dec	.7294	.7295	.7285	.7282	+.0016	.7670	.7270	640
	Est vol 5,389; vol Fri 4,248; open int 44,905, -1,331							

3. What is the CME contract size for:

   (a) Japanese yen?
   (b) New Zealand Dollar?
   (c) Swiss Franc?

4. What is the open interest for the September contract for:
   (a) Japanese yen?
   (b) New Zealand Dollar?
   (c) Swiss Franc?

5. What are the daily high, low, and settlement prices for the December contract for:
   (a) Japanese yen?
   (b) New Zealand Dollar?
   (c) Swiss Franc?

6. What is the day's cash flow from marking to market for the holder of a:
   (a) JPY June contract?
   (b) USD June contract?
   (c) GBP June contract?

7. What statements are correct? If you disagree with one or more of them, please put them right.
   (a) Margin is a payment to the bank to compensate it for taking on credit risk.
   (b) If you hold a forward purchase contract for JPY that you wish to reverse, and the JPY has increased in value, you owe the bank the discounted difference between the current forward rate and the historic forward rate, that is, the market value.
   (c) If the balance in your margin account is not sufficient to cover the losses on your forward contract and you fail to post additional margin, the bank must speculate in order to recover the losses.
   (d) Under the system of daily recontracting, the value of an outstanding forward contract is recomputed every day. If the forward rate for GBP/NZD drops each day for ten days until the forward contract expires, the purchaser of NZD forward must pay the forward seller of NZD the market value of the contract for each of those ten days. If the purchaser cannot pay, the bank seizes his or her margin.

6.7.2 Applications

1. Innovative Bicycle Makers of Exeter, UK, must hedge an accounts payable of MYR 100,000 due in 90 days for bike tires purchased in Malaysia. Suppose that the GBP/MYR forward rates and the GBP effective returns are as follows:

Time	t=0	t=1	t=2	t=3
Forward rate	4	4.2	3.9	4
Effective return	12%	8.5%	4%	0%

   (a) What are IBM's cash flows given a variable-collateral margin account?

   (b) What are IBM's cash flows given periodic contracting?

2. On the morning of Monday, August 21, you purchased a futures contract for 1 unit of CHF at a rate of USD/CHF 0.7. The subsequent settlement prices are shown in the table below.

   (a) What are the daily cash flows from marking to market?

   (b) What is the cumulative total cash flow from marking to market (ignoring discounting)?

   (c) Is the total cash flow greater than, less than, or equal to the difference between the price of your original futures contract and the price of the same futures contract on August 30?

August	21	22	23	24	25	28	29	30
Futures rate	0.71	0.70	0.72	0.71	0.69	0.68	0.66	0.63

3. On November 15, you sold ten futures contracts for 100,000 CAD each at a rate of USD/CAD 0.75. The subsequent settlement prices are shown in the table below.

   (a) What are the daily cash flows from marking to market?

   (b) What is the total cash flow from marking to market (ignoring discounting)?

   (c) If you deposit USD 75,000 into your margin account, and your broker requires USD 50,000 as maintenance margin, when will you receive a margin call and how much will you have to deposit?

November	16	17	18	19	20	23	24	25
Futures rate	0.74	0.73	0.74	0.76	0.77	0.78	0.79	0.81

4. On the morning of December 6, you purchased a futures contract for one EUR at a rate of INR/EUR 55. The following table gives the subsequent settlement prices and the p.a. bid-ask interest rates on a INR investment made until December 10.

   (a) What are the daily cash flows from marking to market?

   (b) What is the total cash flow from marking to market (ignoring discounting)?

   (c) If you must finance your losses and invest your gains from marking to market, what is the value of the total cash flows on December 10?

December	6	7	8	9	10
Futures price	56	57	54	52	55
Bid-ask interest rates, INR, % *p.a.*	12.00-12.25	11.50-11.75	13.00-13.25	13.50-13.75	NA

5. You want to hedge the EUR value of a CAD 1m inflow using futures contracts. On Germany's exchange, there is a futures contract for USD 100,000 at EUR/USD 1.5.

   (a) Your assistant runs a bunch of regressions:

      i. $\Delta S[\text{EUR/CAD}] = \alpha_1 + \beta_1\ \Delta f[\text{USD/EUR}]$

      ii. $\Delta S[\text{EUR/CAD}] = \alpha_2 + \beta_2\ \Delta f[\text{EUR/USD}]$

      iii. $\Delta S[\text{CAD/EUR}] = \alpha_3 + \beta_3\ \Delta f[\text{EUR/USD}]$

      iv. $\Delta S[\text{CAD/EUR}] = \alpha_4 + \beta_4\ \Delta f[\text{USD/EUR}]$

      Which regression is relevant to you?

   (b) If the relevant β were 0.83, how many contracts do you buy? sell?

6. In the previous question, we assumed that there was a USD futures contract in Germany, with a fixed number of USD (100,000 units) and a variable EUR/USD price. What if there is no German futures exchange? Then you would have to go to a US exchange, where the number of EUR per contract is fixed (at, say, 125,000), rather than the number of USD. How many USD/EUR contracts will you buy?

7. A German exporter wants to hedge an outflow of NZD 1m. She decides to hedge the risk with a EUR/USD contract and a EUR/AUD contract. The regression output is, with t-statistics in parentheses, and $R^2 = 0.59$:

$$\Delta S_{[\text{EUR/NZD}]} = a + \underset{(1.57)}{0.15\,\Delta f_{[\text{EUR/USD}]}} + \underset{(17.2)}{0.7\,\Delta f_{[\text{EUR/AUD}]}}$$

   (a) How will you hedge if you use both contracts, and if a USD contract is for USD 50,000 and the AUD contract for AUD 75,000?

(b) Should you use the USD contract, in view of the low t-statistic? Or should you only use the AUD contract?

Chapter 7

Markets for Currency Swaps

As already discussed in Chapter 5, the choice of the currency of borrowing may be difficult; for instance, the currency that offers the lowest PV-ed risk spread may not be the most attractive one from the risk-management point of view. We also know how a firm can nevertheless have its cake and eat it: one can borrow in the low-cost denomination, and then swap that loan into the desired currency. The case we looked at was the simplest possible loan—one with just one single future payment, standing for interest and principal. In that case we (*i*) convert the upfront inflow via a spot transaction from the currency of borrowing into the desired one, and (*ii*) convert the future outflow in the forward market, thus again replacing the currency of the loan's original outflow by the desired one.

But most loans with a life exceeding one year are, of course, multi-payment: interest is typically due at least once a year, and often even twice or four times; and also the principal can be amortized gradually rather than in one shot at the end. To swap such a loan, one would need as many forward hedges as there are future payments. The modern currency swap provides an answer to this: in one contract the two parties agree upon not just the spot conversion in one direction, but also the reverse conversions for all future service payments. The contract is typically set up such that the time pattern of the final payments corresponds to the time pattern of the original. For example, if the original loan is a bullet loan with a fixed interest rate, then the swapped package can also be of the bullet type and with a constant coupon. That last feature would not be achievable with a set of forward contracts: if the original loan has a constant coupon, then the converted coupons will vary depending on their due dates because the forward rates that we use for the conversion depend on the due date. With modern swaps we can even transform a currency-A *floating-rate* loan into a fixed-rate loan in currency B, something which cannot be done with simple forward contracts since the future service payments are not even known yet. So modern swaps are a general and flexible device to change one loan, chosen perhaps because of its low cost, into another loan that for some reason is viewed as more desirable. The second loan could be different from the original one in terms of currency, or interest payments (fixed versus floating), or both. And these are just the plain-vanilla cases; many *ad-hoc* structures can be arranged at the customer's request.

This chapter is structured as follows. In the first section, we consider a landmark deal between two highly respected companies, the currency swap between IBM and the World Bank negotiated in 1981 (commonly viewed as the mother of the modern swaps) and we indicate the subsequent evolution of the swap into a standard, off-the-shelf product. We then show, in Section 7.2, how the modern currency swap works, and why such deals exist. An even more popular variant of currency swap is the interest-rate swap or fixed-for-floating rate swap, which we discuss in Section 7.3. Section 7.4, discusses a combination of the currency swap and the interest swap, called the fixed-for-floating currency swap or circus swap. Section 7.5 concludes this chapter.

7.1 How the Modern Swap came About

From Chapter 5 we know how spot-forward swaps can be used to transform one zero-coupon loan into a zero-coupon loan in a different currency. Swaps can also be used in themselves, as a package of back-to-back loans. The problem is that many of the applications are somewhat shady: shirking taxes, avoiding currency controls, not to mention laundering money. For this reason, back-to-back and parallel loans or spot-forward swaps were for a long time viewed as not 100% respectable. In 1981 all that changed. Two quite-above-board companies, IBM and the World bank, set up a contract which was quite clever and had a respectable economic purpose: avoiding transaction costs. There was a tax advantage too, but this was almost by accident.

The IBM-WB swap was a bilateral deal, very much tailor-made. But rapidly the swap became a standardised product offered routinely by banks. This evolution is depicted after our description of the IBM-WB deal.

7.1.1 The Grandfather Tailor-made Swap: IBM-WB

In 1981, IBM wanted to get rid of its outstanding DEM- and CHF-denominated callable debt because the USD had appreciated considerably and the DEM and CHF interest rates had also gone up. As a result of these two changes, the market value of IBM's foreign debt, expressed in terms of DEM and CHF, was below its face value, and the gap between market value and book value was even wider in terms of USD. IBM wanted to lock in this capital gain by replacing the DEM and CHF debt by new USD debt. However, in order to do this, IBM would have to incur many costs:

- IBM would have to buy DEM and CHF currency, thus incurring transaction costs in the spot market. In 1981 this was not yet the puny item it has become by now.

- Much more importantly, IBM's loans were callable indeed (that is, IBM could amortize them early)—but at a price above par. So IBM would have to fork out more than the DEM and CHF face value rather the economic value of the straight-bond component, which was below par. Calling would be like exercising an out-of-the-money option.

  Finance theorists will happily point out that hedging the debt would be the obvious solution: borrow dollars and invest them in DEM and CHF assets that match the outstanding debt, thus neutralizing any possible re-appreciation of these currencies without any need to actually withdraw the old bonds. But CFO's will unhappily note that, in conventional accounting terms, this would double the debt.

- IBM would have to pay a capital-gains tax on the difference between the (dollar) book value and the price it paid to redeem the bonds.

- Lastly, IBM would have to issue new USD bonds to finance the redemption of its CHF and DEM debt. In those days, a bond issue costed at least a few percentages of the nominal value.

The World Bank (WB), on the other hand, wanted to borrow DEM and CHF to lend to its customers. Its charter indeed said that, currency by currency, its assets should be matched by its liabilities. Clearly, issuing new CHF and DEM bonds would have entailed issuing costs.

To sum up, IBM wants to *withdraw* CHF and DEM bonds (at a rather high cost) while WB wants to *issue* CHF and DEM bonds (also at a cost) (see Figure 7.1). To avoid all of these expenses, IBM and WB agreed that

- the WB would not borrow CHF and DEM, but would borrow USD instead. With the proceeds, it would buy spot CHF and DEM needed to make loans to its customers.

- the WB undertook to take over the servicing of IBM's outstanding DEM and CHF loans, while IBM promised to service the WB's (new) USD loan.

This way, each party achieved its objective. IBM has effectively traded (or swapped) its DEM and CHF obligations for USD obligations: its DEM-CHF debt is taken care of by WB, economically, and IBM now services USD bonds. The WB, on the other hand, has an obligation to deliver DEM and CHF, which is what WB needed. One obvious joint saving of the swap was the cost of issuing new WB bonds in DEM and CHF, and redeeming the old IBM loans in DEM and CHF. Also, the recognition of IBM's capital gain was postponed because the old bonds were not redeemed early. Another saving was that the WB could issue USD bonds at a lower risk spread than IBM.[1]

[1]Critical economist would rightly object that this is a saving only to the extent the difference in the risk spreads was irrational.

Figure 7.1: **The IBM-WB swap**

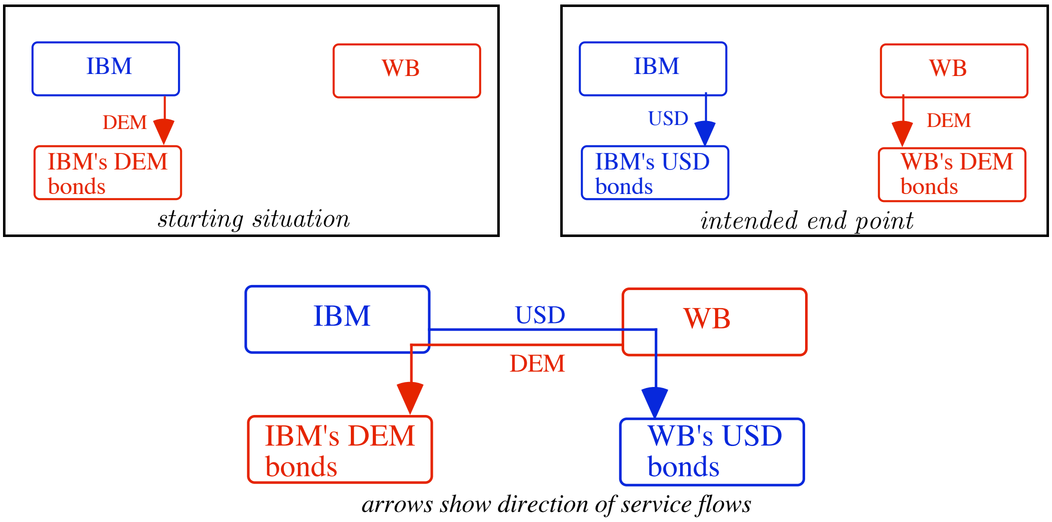

Key Top left: the initial situation; top-right: the originally intended final situation; bottom: how the essence of the desired solution was realized, at a lower cost.

.

Of course, the amounts to be exchanged had to be acceptable to both parties. The present value of IBM's USD payments to the WB should, therefore, be equal to the present value of the DEM and CHF inflows received from the WB.

Example 7.1

Assume, for simplicity, that IBM has an outstanding DEM debt with a face value of DEM 100m and a book value of USD 60m (based on the historic USD/DEM rate of 0.6), maturing after five years and carrying a 5 percent annual coupon. Assume the current five-year DEM interest rate is 10 percent and the DEM now trades at USD/DEM 0.4. In DEM, IBM's existing debt would have a present value of[2]

$$\text{DEM}\,100m \times [1 + (0.05 - 0.1) \times a(10\%, 5 \text{ years})] = \text{DEM}\,81.05m, \tag{7.1}$$

where $a(r, n)$ is the present value of an n-year unit annuity discounted at a rate r:

$$a(r, n) \stackrel{\text{def}}{=} \sum_{t=1}^{n} \frac{1}{(1+r)^t} = \frac{1 - (1+r)^{-n}}{r}. \tag{7.2}$$

At the current spot rate of USD/DEM 0.4, WB's undertaking to service this debt is worth $81.05 \times 0.4 =$ USD 32.42m.

[2]See TekNote 7.1 if the formula is new to you.

The equal-value principle requires that IBM's undertaking have the same present value. Thus, the USD loan (issued at the then-prevailing rate for five years) must have a present value of USD 32.42m.

As we have argued, one purpose of the entire IBM/WB deal was to avoid transaction costs. A nice by-product, in terms of taxes, was that IBM locked in its capital gain on its foreign currency debt without immediately realizing the profit. Let us quantify some of these elements using the above figures. If IBM had called its DEM debt at 102 percent of its DEM par value, the cost of withdrawing the debt would have been 100m $\times$ 1.02 $\times$ USD/DEM 0.4 = USD 40.8m, thus realizing a taxable capital gain of USD 60m – 40.8 = USD 19.2m. In contrast, under the swap, the DEM debt remains in IBM's books for another five years. That is, in accounting terms, the capital gain will be realized only when, five years later, IBM pays the swap principal (USD 32.42m) to the WB and receives DEM 100m to redeem its DEM debt. In short, the swap also allowed IBM to defer its capital gains taxes.

7.1.2 Subsequent Evolution of the Swap Market

We know that a forward contract is like an exchange of two initially equivalent promissory notes, one in HC and one in FC. In the IBM-WB deal we see, instead, something like an exchange of two bonds (or a least cash-flow patterns that correspond to bond servicing schedules). This differs from the forward contract in the sense that there is not just an exchange of two main amounts at the end, but also interim interest is being paid to each other at regular dates. But the principle of initial equivalence of the two "legs" of the deal is maintained.

One feature that has changed nowadays, relative to the IBM-WB example, is that almost invariably a reverse spot exchange is added. One reason is that very often the purpose of the swap is to transform a loan taken up in currency X into one expressed in currency Y; and to do that, one also needs the immediate currency-Y inflow beside the future outflows.

DoItYourself problem 7.1

Suppose you want to borrow GBP, but what you actually do is borrow USD and swap, the way we saw it in Chapter 5. So part of the deal is that you promise the swap dealer a stream of GBP; the swap dealer in return then pays you USD with which you can service your bank loan. But all this only delivers you the future-GBP-outflow part of the desired loan. To get also the immediate-GBP-inflow part, you convert the USD proceeds of the bank loan into pounds.

A second reason for adding the spot deal is that the exchange of time-t PVs simplifies the negotiation process. One has to realize, indeed, that swaps are typically add-ons to biggish loans; and taking up a big loan is a much rarer and slower decision than, say, a spot or forward transaction that has to do with trade transactions. Since negotiations take hours or days, and since the spot rate is moving all the time, one would have to continuously change one leg of the swap to maintain initial equivalence of the future payments. By throwing in an exchange of the spot PVs, this problem is much reduced. The idea is that one can still get zero initial value for the swap as a whole if the *Net* PV of each leg separately is zero—the PV of the future payments minus the initial flow in the opposite direction.

Example 7.2

Suppose that, at the beginning of the negotiations, a US company promises to send to a Dutch company a stream of USD corresponding to a bullet loan with notional value USD 50m at 4 percent payable annually. Suppose the normal yield rate for this type of bond is 4 percent, so that PV_{USD} = USD 50m. Company B promises a stream of EUR in return. On the basis of S_t = USD/EUR 1.25 and a EUR interest rate of 4.5 percent, the EUR payments would mirror the service payments for a EUR 40m loan at 4.5 percent. This way, the PVs of the EUR and USD streams are identical, resulting in a zero total value of the contract.

But if one hour later the spot rate is 1.26, the calculations would have to be revised. This revision would become unnecessary if the contract also stipulates an initial exchange of EUR 40m for USD 50m. Then, to the US company, the appreciation of the EUR increases the USD PV of the incoming future Euros but also increases by the same factor the USD value of the EUR amount the company needs to fork out immediately. Thus, the net value of the EUR leg remains zero as long as interest rates do not change.

With the immediate exchange of principals brought in, one can do with approximate equivalence of the two notional amounts. An approximate equivalence is still important because the two loans also serve as security for each other. If one side were far smaller than the other, the security provision would be unacceptably asymmetric.

A second major change, relative to the IBM-WB example is that contracts are now standardized. The early swaps were carefully negotiated between two parties, with task forces of financial economists and lawyers in attendance to calculate the gains and to arrive at a fair division of the gains. Inevitably, then, one huge initial problem was to find a counterpart with the complementary objectives. In forward markets, we know, banks act as intermediaries. If company A buys forward, the bank agrees, and afterwards solicits a sale from someone else by skewing its bids (Chapter 3), or the bank closes out in the spot and money markets (synthetic sale). This is exactly how things have become in the swap markets too. A company signs a swap agreement with a bank, which may keep this contract "on its book" (*i.e.* open) for a while, until new contracts have brought the overall book closer to neutrality. If the risk is too large, the bank can always hedge in the bond and spot markets (synthetic swap). This hedging was easiest in the USD interest-rate swap market, where the two notional loans that constitute the swap are expressed in the same currency but have different interest forms—typically, one leg fixed-rate and the other floating-rate. Given that there is a huge market for similar fixed- and floating-rate bonds outstanding, swap dealers could easily close out in the bond market. Also, a lively secondary market for swaps has emerged.

We are now ready to have a closer look at the how swaps are set up. We begin with "fixed-for-fixed" currency swaps, that is, swaps with fixed coupons in each leg.

7.2 The Fixed-for-Fixed Currency Swaps

7.2.1 Motivations for Undertaking a Currency Swap

The reasons for using swaps are essentially those mentioned for spot-forwards swaps (Chapter 5). Generally, the point is to avoid unnecessary costs generated by market imperfections, primarily information costs that lead to excessive risk spreads asked by uninformed banks. The IBM-WB case was mainly a transaction-cost motivated structure. Also the advantages of off-balance-sheet reporting remain valid, at least when a swap is compared to its synthetic version (borrow in one currency and invest in another).

An extreme form of market imperfection arose in one particular instance: in the early 1980s, the French car manufacturer Renault wanted to borrow Yen and use the proceeds to redeem outstanding USD debt, but found that in those days the Yen bond market was quasi closed to foreign borrowers. So Renault swapped its USD loan with Yamaichi Securities for JPY debt. The Renault-Yamaichi swap was not a fixed-for-fixed swap, so its discussion is deferred to Section 7.4, below.

7.2.2 Characteristics of the Modern Currency Swap

In many ways, the modern fixed-for-fixed currency swap is simply a long-term version of the classical spot-forward swap. A fixed-for-fixed currency swap can be defined as a transaction where two parties exchange, at the time of the contract's initiation, two principals denominated in different currencies but with (roughly) the same market value, and return these principals to each other when the contract expires. In addition, they periodically pay a normal interest to each other on the amounts borrowed. The deal is structured as a single contract, with a right of offset. The features of a fixed-for-fixed currency swap are described in more detail below.

Swap rates

In a fixed-for-fixed currency swap, the interest payments for each currency are based on the currency's "swap (interest) rate" for the swap's maturity. These swap rates are simply yields at par for near-riskless bonds with the same maturity as the swap.[3] In practice, the swap rates are close to the long-term offshore rates on high-quality sovereign loans, that is, loans by governments. For the following reasons, it is appropriate to use near-risk-free rates to compute the interest on the amounts swapped even if the counterpart in the contract is not an AAA company:

- The bank's risks in case of default are limited because of the right-of-offset clause. In unusually risky cases, the contract parties also have to post margin.
- The probability of default is small. This is because the customers are screened; small or low-grade companies get no chance, or have to post initial margin.
- In addition, many swap contracts have a "*credit trigger*" clause, stating that, if the customer's credit rating is revised downward, the financial institution can terminate the swap, and settle for the swap's market value at that moment. Thus, the bank has an opportunity to terminate the contract long before default actually occurs—unless the company goes straight from AA to failure, Enron-style.
- Finally, because of the right of offset, the uncertainty about the bank's inflows is the same as the uncertainty about the bank's outflows. The fact that the uncertainties are the same implies that the corrections for risk virtually cancel out. That is, it hardly matters whether or not one adds a similar (and small) default risk premium to the risk-free rates when one discounts the two cash flow streams. The effect of adding a small risk premium when valuing one "leg" of the swap will essentially cancel out against the effect of adding a similar risk premium in the valuation of the other leg.

Look at the rates in Figure 7.2. Sterling has a one-year swap rate of 4.96-4.99. Elsewhere in the same FT issue I find the following 1-year rates: Interbank Sterling 4.875-4.96875, BBA Sterling 4.65625, Sterling CD 4.90625-4.9375, Local authority depts 4.875-4.9375. Thus, the swap rate is close to a risk-free rate. There *is* a small risk premium, but it is so low that for all practical purposes you can think of the swap rate as the risk-free rate, the same way LIBOR is called risk-free.

The Key to the FT table mentions another detail: a swap rate is quoted against a particular floating rate. This is from the fact that the busiest section of the swap market is the interest swap, fixed versus floating or vice versa. In principle it should not matter what exactly the floating-rate part is: since investors can freely

[3] The N-year yield at par is the coupon that has to be assigned to an N-year bond in order to give it a market value equal to the par value (the principal). If the parties want a cash flow pattern that differs from the single-amortization ("bullet") loan, the swap dealer is usually willing to design a contract that deviates from the standard form, but at a different swap rate. See the subsection on non-bullet loans, further in this section.

Figure 7.2: **Swap Rates as quoted in the Financial Times**

	Euro-C		£ Stlg		SwFr		US $		Yen	
06/06/06	bid	ask	bid	ask	bid	ask	bid	ask	bid	ask
1 year	3.42	3.45	4.96	4.99	1.91	1.97	5.46	5.49	0.60	0.62
2 year	3.64	3.67	5.03	5.07	2.26	2.34	5.41	5.43	0.92	0.95
3 year	3.77	3.80	5.08	5.12	2.46	2.54	5.41	5.43	1.18	1.21
4 year	3.86	3.89	5.09	5.14	2.59	2.67	5.42	5.45	1.39	1.42
5 year	3.93	3.96	5.09	5.14	2.69	2.77	5.45	5.47	1.56	1.60
6 year	3.99	4.02	5.09	5.14	2.78	2.86	5.46	5.50	1.71	1.74
7 year	4.05	4.08	5.08	5.12	2.85	2.93	5.49	5.51	1.84	1.87
8 year	4.10	4.13	5.06	5.11	2.91	2.99	5.51	5.54	1.95	1.98
9 year	4.15	4.18	5.04	5.09	2.97	3.05	5.53	5.56	2.04	2.06
10 year	4.20	4.23	5.01	5.07	3.02	3.10	5.55	5.58	2.11	2.14
12 year	4.28	4.31	4.96	5.03	3.08	3.18	5.58	5.62	2.24	2.27
15 year	4.38	4.41	4.88	4.97	3.17	3.27	5.63	5.66	2.38	2.41
20 year	4.47	4.50	4.75	4.88	3.26	3.36	5.66	5.69	2.54	2.57
25 year	4.51	4.54	4.64	4.77	3.27	3.37	5.66	5.70	2.63	2.66
30 year	4.52	4.55	4.56	4.69	3.26	3.36	5.66	5.69	2.67	2.70

Bid and ask rates are as of close of London business. US $ is quoted annual money actual/360 basis against 3 months Libor, pound and Yen quoted on a semi-annual actual/365 basis against 6-months Libor, Euro/Swiss Franc rate quoted on annual bond 30/360 basis against 6-month Euribor Libor with the exception of the 1 year rate which is quoted against 3 month Euribor/Libor. Source: ICAP plc

chose between say 3- and 6- or 9-month LIBOR, the three should be equivalent. In practice, differences in e.g. liquidity may cause the swap rate to differ, in a minor way, depending on what the floating-rate part is.

Costs

The swapping bank charges a small annual commission of, say, USD 200 on a USD 1m swap, for each payment to be made. Most often this fee is built into the interest rates, which would raise or lower the quoted rate by a few basis points.

Example 7.3

Suppose that the seven-year yields at par are 3.17 percent on USD and 3.9 percent on EUR. The swap dealer quotes

USD 3.15%–3.19%,
EUR 3.88%–3.92%.

If your swap contract is one where you "borrow" EUR and "lend" USD, you would then pay 3.92 percent on the EUR, and receive 3.15 percent on the USD.

Theoretically, the series of future commissions, one per payment, might be replaced by a single up-front fee with a comparable present value. Even if this is seldom done in practice, it is still useful for you to always calculate this number, so as to have an idea of the overall cost. For a ten-year USD 1m swap at 3 percent annually that has a USD 200 commission per payment, the equivalent up-front commission would be about 200 × a(3%, 10 years) = 200 × 8.530,203 = USD 1706, or

Table 7.1: **Fixed-for-fixed Currency Swap: the Interim Solution**

	loan	swap		Combined
	JPY 1000 borr'd at 1%	JPY 1000 lent, at 0.6%	USD 10m borr'd at 3%	
principal at t	JPY 1000m	<JPY 1000m>	USD 10m	USD 10m
interest (p.a.)	<JPY 10m>	JPY 6m	<USD 0.3m>	<JPY 4m>& <USD 0.3m>
principal at T	<JPY 1000m>	JPY 1000m	<USD 10m>	<USD 10m>

0.17 percent of the face value.

Thus, although the swap remains a zero-value contract, the customer has to pay a small commission. (You can tell the difference between a price and a commission because the commission is always paid, whether one goes long or short; in contrast, the price is paid if one buys, and is received if one sells.) The commissions in the swap are small because the costs of bonding and monitoring are low—default risk is minimal anyway, as we saw—and because the amounts are large. (A typical interbank swap transaction is for a few million USD, and the Reuters swap-dealing network requires minimally USD 10m; for corporations, swaps can be smaller but contracts below USD 1m are rare.) Familiarly, the swap spread also depends on liquidity. Deep markets like USD, EUR and JPY, in Figure 7.2, have spreads of 3 bp or thereabouts, but for CHF and GBP the margin is wider, rising to 10 bp at the far end of the maturity spectrum.

How to Handle and Compare Risk Spreads

Suppose a Japanese company wants to borrow cheaply in JPY (=HC) from its house bank, at 1 percent for 7 years, bullet, and then swap the loan into USD. The swap rates quoted are 0.6 percent on JPY and 3 percent on USD. In Table 7.2 I set up a little tabular that shows you, in the first column of figures, the original loan (JPY at 1%); in the next two, the twin legs of the swap; and, lastly, the combined cash flow (loan and swap). The version I show in that table is, actually, rarely applied in practice; I mainly use it as an interim step because it helps to explain the advantage of the swap as well as the logic of the ultimate solution. The spot rate being about JPY/USD 100, we work with notional principals of JPY 1000m and USD 10m. So the company

– borrows JPY 1000m from the house bank at 1 percent (the actual loan rate),

– "re-lends" these JPY 1000m to the swap dealer, at 0.6% (the JPY swap rate),

– ... who in return "lends" USD 10m to the firm at 3% (the USD swap rate).

This is summarized in Table 7.1. Note how the company borrows, ultimately,

Table 7.2: **Fixed-for-fixed Currency Swap: Marked-up USD rate**

	loan	swap		Combined
	JPY 1000 borr'd at 1%	JPY 1000 lent, at 1%	USD 10m borr'd at 3.438 823%	
principal at t	JPY 1000m	<JPY 1000m>	USD 10m	USD 10m
interest (p.a.)	<JPY 10m>	JPY 10m	<USD 343,882.30>	<USD 343,882.30>
principal at T	<JPY 1000m>	JPY 1000m	<USD 10m>	<USD 10m>

USD 10m, with an annual interest payment consisting of the USD risk-free rate (3 percent) plus a risk spread which is, very literally, the risk spread on a JPY loan from the house bank: 1% − 0.6% = 0.4% on JPY 1000m.[4]

The above solution is still somewhat unelegant because the company pays part of its annual interest payments in JPY, an undesirable feature if it basically wants a USD loan. There are two simple solutions:

- either replace the seven annual JPY 4m payments by an equivalent upfront fee, which is of course their PV:

$$\begin{aligned}\text{Equivalent upfront fee} &= 4m \times a(0.6\%, 7years) \\ &= 4m \times 6.834,979 = \text{JPY}\,27.339,917m, \quad (7.3)\end{aligned}$$

- or replace it by an equivalent USD annuity:

$$\text{Find } X^* \text{ such that} \quad \underbrace{S_t \times \overbrace{X^* \times a(3\%, 7years)}^{\text{PV of annuity USD } X^*}}_{\text{... translated into JPY}} = \underbrace{4m \times a(0.6\%, 7years)}_{\text{PV of annuity JPY 4m}}$$

$$\begin{aligned}\Rightarrow X^* &= \frac{4m}{S_t}\frac{a(0.6\%, 7years)}{a(3\%, 7years)} \qquad (7.4)\\ &= \frac{4m}{100}\frac{6.834,979}{6.230,282} = \text{USD}\,43,882.30m\end{aligned}$$

Technically, we ask the swap dealer to pay us 1% (our borrowing rate) on the yens instead of 0.6% (the swap rate), and in return we increase the dollar interest paid to the swap dealer by the equivalent amount. Table 7.2 summarizes the modified solution.

The second solution immediately allows us to discover whether the swapped loan is more attractive than a direct USD loan (an alternative we have not yet looked

[4] Note that, since these two amounts are in different currencies with a stochastic future exchange rate, there is no way to amalgamate them into one number or one percentage. That is, 3 percent in USD and 0.4 percent in JPY is not 3.4 percent.

at). The translated risk spread equivalent to the 0.4% charged by the Japanese housebank, as a percentage of the USD 10m borrowed, amounts to 43,882.30/10m = 0.438823 %. Let's denote the risk spreads by ρ and ρ^*, as in Chapter 5, and let's use s and s^* to refer to the swap rates. You can check that the generalized equivalence condition is

$$\rho^* \overset{equiv}{=} \rho \frac{a(s,n)}{a(s^*,n)} \Leftrightarrow \underbrace{\rho^* \times a(s^*,n)}_{\text{PV of FC risk spread}} \overset{equiv}{=} \underbrace{\rho \times a(s,n)}_{\text{PV of HC risk spread}} \tag{7.5}$$

Thus, a borrower gains from the swap if the spread quoted for a direct loan is higher than this translated HC risk spread, the HC figure projected into a different interest-rate environment via the adjustment $\times a(s,n)/a(s^*,n)$. Similarly, a credit analyst working for a bank can use the formula to consistently translate the borrower's HC risk spread into FC. The solution is a straightforward generalization of the one for simple spot-forward swaps in Chapter 5: a FC risk spread ρ^* is equivalent to a HC ρ if their PVs are the same. The only change is that, of course, the PV'ing now involves annuities rather than a single payment: in a bullet loan, the risk premium is paid many times, not just once. Note also that for non-bullet loans the above formula no longer works, because the risk-spread payments (in amounts, not percentages) then no longer are constant. The equal-PV rule for equivalence still would hold, but the computations would be messier.

Also the intuition as to why and when a translated risk spread exceeds the original one remains the same as before. In risk-adjusted terms, the Yen is the strong currency here, as we can infer from its lower interest rate. So a strip of 0.4 percent payments in USD cannot be as good as a series of 0.4 percent in Yen, the strong currency. The above formula tells us exactly how the strength of the currencies, as embodied in their interest rates, has to be quantified in the translation process: taking into account the relative annuity factors, one needs to offer 0.438823 % in USD to be in balance with 0.4% in Yen.

Non-bullet loans

Standard swap-rate quotes are for bullet loans. Any other package is replicated as a combination of bullet loans with different times to maturity, and for each component the appropriate swap rate holds.

Example 7.4

Assume the swap rates for 1, 2, and 3 years are 5, 6, and 7%, respectively. We want to create a three-year constant-annuity loan, with three payments worth 1000 each. The tools we have are three bullet loans: a one-year specimen with face value V_1 (to be determined); a two-year one with face value V_2; and a three-year loan with face value V_3.

Finding the replication requires solving a simple linear system. In the case of a

Table 7.3: **Replicating a constant-annuity loan from bullet loans**

	interest payments on loan maturing in year ...			amortization payments on loan maturing in year ...			
	1	2	3	1	2	3	sum
year 1	41.985	52.901	65.421	839.694	0	0	1000
year 2	0	52.901	65.421	0	881.679	0	1000
year 3	0	0	65.421	0	0	934.579	1000

Key The loans are 839.694 for one year, 881.679 for two years, and 934.579 for three — just believe me, or read Figure 7.3. The annual interest payments are 5% (one year loan), 6% (two) or 7% (three), and each loan is amortized on the promised dates. The total combined service schedule is exactly 1000 every year.

Figure 7.3: **Replicating a constant-annuity loan from bullet loans**

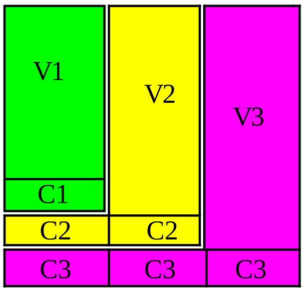

$V_3 + C_3 = 1000$
$\mathbf{V_3}\,(1+s_3) = 1000 \Rightarrow V_3$

$V_2 + C_2 + C_3 = 1000$
$\mathbf{V_2}\,(1+s_2) + V_3\,s_3 = 1000 \Rightarrow V_2$

$V_1 + C_1 + C_2 + C_3 = 1000$
$\mathbf{V_1}\,(1+s_1) + V_2\,s_2 + V_3\,s_3 = 1000 \Rightarrow V_1$

V_3 = 934.58 V2 = 881.68, V1 = 839.69

Key A service schedule amounting to three times 1000 is arranged as follows:

- We begin with year 3. In that year, only the three-year loan is still alive, and its total service cost including 7% interest must be 1000. So the requirement is to find V_3 such that $V_3 \times 1.07 = 1000$—that is, $V_3 = 1000/1.07 = 934.579$. The balance is interest on the 3-year bullet loan.
- Of the total 1000 paid in year 2, the same 934.579 is available, after paying the interest in the three-year bullet loan, for principal and coupon of the two-year bullet loan: $V_3 = V_2 \times 1.06$. So $V_2 = V_3/1.06 = 881.678$, etc.

constant-annuity loan the rule is that $V_t = V_{t+1}/(1+s_t)$, with a "dummy" V_4 defined as the annuity itself—that is, $V_3 = 1000/1.07 = 934.579$, $V_2 = 934.579/1.06 = 881.678$, and $V_1 = 881.679/1.05 = 839.694$. Table 7.3 verifies that this indeed produces a combined cash flow for the three loans together of 1000 every year, and Figure 7.3 show you how you get these numbers.

This way, the swap dealer has also found that the PV of the three-year annuity is 934.58 + 881.68 + 839.69 = 2,655.95. Spreadsheet afficionados will readily confirm that this corresponds to an IRR of 6.347 %. This would then be the swap dealer's rate for three-year constant-annuity loans. As you see, this is neither the 3-year rate nor the 2- or 1-year rate for bullet loans, but a complicated mixture of all three. But this is the swap dealer's problem: the user can just work with the swap rate given to her for the particular type of loan at hand.

Valuing an Outstanding Fixed-for-Fixed Currency Swap

The last issue that we discuss in the section on fixed-for-fixed currency swaps is the valuation of such a swap after its inception. An assessment of the market value of a swap is required for the purpose of true and fair reporting to shareholders and overseeing authorities, or when the contract is terminated prematurely (by negotiation, or by default, or by the credit trigger clause).

Just as a forward contract, the fixed-for-fixed currency swap acquires a non-zero value as soon as the interest rates change, or as the spot rate changes. Since a swap is like a portfolio of (a) a loan and (b) an investment in long-term deposits (or in bonds), we can always value a swap as the difference between the market value of the loan and the market value of the investment.

Example 7.5

Two years ago, a bank swapped a company loan (asset) of USD 100m for GBP 50m for seven years, at the swap rates of 4 percent on the USD leg and 5 percent on the GBP leg. This reflected the long-term interest rates and the spot rate of USD/GBP 2 prevailing when the contract was signed. Now the five-year USD swap rate is 2.5 percent, the five-year GBP swap rate is 4 percent, and the spot rate is USD/GBP 1.7. The procedure suggested by the International Swap Dealers Association is to value the swap by applying the traditional bond valuation formula to each of the swap's legs. Thus, the company's USD outflows are valued as

$$\text{PV}_{\text{USD}} = 100m \times [1 + (0.04 - 0.025) \times a(2.5\%, 5years)] = \text{USD}\,106,968,742.74, \quad (7.6)$$

while its GBP inflows are worth

$$\text{PV}_{\text{GBP}} = 50m \times [1 + (0.05 - 0.04) \times a(4\%, 5years)] = \text{GBP}\,52,225,911.17. \quad (7.7)$$

At the spot rate of USD/GBP 1.7, these GBP inflows are worth USD 88,784,048.98. The contract has therefore become a net liability, with value USD 88,784,048.98 – 106,968,742.74 = –USD 18,184,693.76.

This finishes our discussion of the fixed-for-fixed currency swap. We now turn to other types of swaps, the most important of which is the interest rate swap or coupon swap.

7.3 Interest Rate Swaps

In an interest rate swap, there is still an exchange of the service payments on two distinct loans. However, the two loans involved now differ not by currency, but by the method used to determine the interest payment (for instance, floating rate versus fixed rate). Because both underlying loans are in the same currency, there is no initial exchange of principals and no final amortization. In that sense, the two loans are *notional* (fictitious, or theoretical). The only cash flows that are swapped are the interest streams on each of the notional loans. In short, parties A and B simply agree to pay/receive the difference between two interest streams on the notional loan amounts.

The standard interest swap is the fixed-for-floating swap or coupon swap. The base swap is rarer. We discuss each of them in turn.

7.3.1 Coupon Swaps (Fixed-for-Floating)

We now describe the characteristics of a fixed-for-floating swap and how one can value such a financial contract.

Characteristics of the Fixed-for-Floating Swap

In our discussion of the fixed-for-fixed currency swap, we saw that, in terms of the risk spread above the risk-free rate, a firm often has a comparative advantage in one currency but may prefer to borrow in another currency. The firm can retain its favorable risk spread and still change the loan's currency of denomination by borrowing in the most favorable market and swapping the loan into the preferred currency. The same holds for the fixed-for-floating swap except that, instead of a preferred currency, the firm now has a preferred type of interest payment. For instance, the firm may have a preference for financing at a fixed rate, but the risk spread in the floating-rate market may be lower. To retain its advantage of a lower spread in the floating-rate market, the firm can borrow at a floating rate, and swap the loan into a fixed-rate loan using a fixed-for-floating swap.

Because the swap contract is almost risk free, the interest rates used in the swap contract are (near) risk-free rates. For the floating-rate leg of the swap, the rate is traditionally LIBOR or a similar money market rate, while the relevant interest rate for the fixed-rate leg is the same N-year swap rate as used in fixed-for-fixed currency swaps. In fact, traditionally, the fixed swap rate was *defined* as the rate which the swap dealer thought to be as good as LIBOR, that is, which she or he was willing to take as the fixed-rate leg in a fixed-for-floating or floating-for-fixed swap. Also, LIBOR in currency X is also defined as acceptable against LIBOR in currency Y, which in turn must be acceptable against currency-Y fixed.

Example 7.6

An AA Irish company wants to borrow NZD to finance (and partially hedge) its direct investment in New Zealand. Because the company is better known in London than in Auckland, it decides to tap the euro-NZD market rather than the loan market in New Zealand. As NZD interest rates are rather volatile, the company prefers fixed-rate

Table 7.4: **Fixed-for-floating swap**

	loan	swap		Combined
	NZD 1 borr'd at LIBOR +1%	NZD 1 lent, at LIBOR %	NZD 1m borr'd at 5%	
principal at t	NZD 1	<NZD 1>	NZD 1	NZD 1
interest (p.a.)	< LIBOR +1%>	LIBOR %	< 5% >	$< 5\% + 1\% = 6\% >$
principal at T	<NZD 1>	NZD 1	<NZD 1>	<NZD 1>

loans. But eurobanks, which are funded on a very short-term basis, dislike fixed-rate loans, which means that the company would have to tap the bond market. The company's alternatives are the following:

- A euro-NZD fixed-rate bond issue would be possible only at 7 percent, which represents a hefty 2 percent spread above the NZD swap rate of 5 percent.
- From a London bank, the Irish company can get a NZD floating-rate bank loan at LIBOR + 1 percent.

The company can keep the lower spread required in the floating-rate market and still pay a fixed rate, by borrowing NZD at the NZD LIBOR + 1 percent, and swapping this into a fixed-rate NZD loan at the 5 percent swap rate. The payment streams, per NZD, are summarized in Table 7.4. To help you see the link between the payments under the swap contract and the underlying notional loans, we have added the theoretical principals at initiation and at maturity. In practice, the principals will not be exchanged. We see that this company borrows foreign currency at the NZD risk-free fixed rate (5 percent) plus the spread of 1 percent it can obtain in the "best" market (the floating-rate eurobank-loan market). Therefore, the company pays 6 percent fixed rather than the 7 percent that would have been required in the bond market.

Having done the number-crunching, let's talk economics now: how it is possible that the bond market requires 2 percent, by way of risk spread, when banks are happy with 1 percent? One reason is that banks are quite good at credit analysis, while Swiss dentists—still a non-trivial part of the bond-market clientèle—are not trained analysts. Also, the amounts at stake for a bank do justify a thorough analysis, while the 10,000 dollars invested by the Swiss dentists are too small for this. Furthermore, our Irish company will be happy to privately provide information to its bank that it would not dream of publishing in a prospectus. In short, the bank knows more, and knows better what the information means.

The swap dealer, who has to find a new party with (roughly) the opposite wants as our AA company, might then talk to an institutional investor, like an insurance company. They like long-duration deals. So everybody is happy. The insurance company gets a long-run fixed-rate investment and the firm the long-run fixed-rate funding, but the credit analysis and the first-line default risk are left to the credit specialist, the bank.

From the above discussion, it is obvious that the potential advantages of the coupon swap are similar to the ones mentioned in the case of the fixed-for-fixed currency swap. What remains to be discussed is how to determine the value of the fixed-for-floating swap.

Valuing a Fixed-for-Floating Swap

We have seen that in a fixed-for-floating swap without default risk, the incoming stream is the service schedule of a risk-free floating-rate loan, and the outgoing stream is the service schedule of a traditional risk-free fixed-rate loan (and vice versa for the other contract party). The fixed-rate payment stream is easily valued by discounting the known cash flows using the prevailing swap rate for the remaining time to maturity. The question now is how should one value the floating-rate part for which the future payments are not known in advance.

Let us study the value of a series of floating-rate cash inflows. This series of (as yet unknown) inflows must have the same market value as a short-term deposit where the principal amount is reinvested periodically. The reason for this equivalence is that the cash flows are the same, as the example will show. To buy such a series of deposits we need to buy only the currently outstanding deposit. No extra money is needed to redeposit the maturing principals later on.

Example 7.7

Suppose that you want to replicate a risk-free USD 10,000 floating-rate bond with semiannual interest payments equal to the 6-month T-bill rate, the first of which is due within four months. At the last reset date, the six-month T-bill rate was 3 percent *p.a.*; thus, the next interest payment equals 10,000 $\times$ (1/2) $\times$ 3% = USD 150. The current four-month rate of return is 0.9 percent (or 2.7 percent *p.a.*, simple interest).

The above floating-rate bond can be replicated by "buying" USD 10,150 due three months from now at a present value cost of USD 10/1.009 = USD 10059.46. After four months, you withdraw USD 150 to replicate the bond's first coupon, and you redeposit the remaining 10,000 at the then prevailing six-month return. When this investment expires, you again withdraw the interest and redeposit the 10,000 at the then-prevailing rate, and you continue to do so until the bond expires. Notice that the future payoffs of the rolled-over deposit are identical to the payoffs of the floating rate bond. The cost to you was only the initial USD 10059.46. Then, by arbitrage, the floating rate bond is also worth 10059.46.

Figure 7.4: **Crucial dates for valuing a floating-rate note**

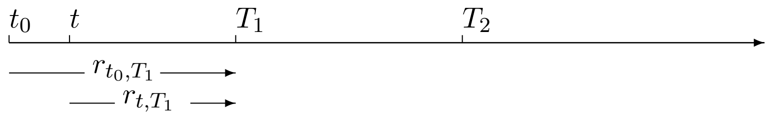

The general expression for the value of a floating-rate bond is

$$\text{Value of a risk-free floating-rate bond } = [\text{Face value}] \times \frac{1 + r_{t_0,T_1}}{1 + r_{t,T_1}}, \tag{7.8}$$

where

t_0	is the last reset date
T_1	is the next reset date
t	is the present date (the valuation date, with $t_0 < t < T_1$)
r_{t_0,T_1}	is the coupon effectively payable at T_1
r_{t,T_1}	is the effective current return until time T_1.

The current market value of a coupon swap then equals the difference between the market value of the loan that underlies the incoming stream and the market value of the loan that underlies the outgoing stream.

Example 7.8

Some time ago, a South-African company speculated on a drop in fixed-rate interest rates, and swapped ZAR 10m at 7 percent, semi-annual and fixed, for ZAR 10m at the six-month ZAR LIBOR. That is, the contract stipulates that, every six months, the firm pays the six-month ZAR LIBOR rate (divided by two)[5] on a notional ZAR 10m, and receives 7/2=3.5 percent on the same notional amount. Suppose that ZAR medium-term interest rates have fallen substantially below 7 percent. The company reckons it has made a nice profit on its swap contract, and wants to lock in the gain by selling the swap. Current conditions are as follows:

- The swap has five years and two months left until maturity.
- The current five-year ZAR swap rate (for semiannual payments) is 5 percent *p.a.*, meaning 2.5 percent every six months.
- The ZAR LIBOR rate, set four months ago for the current six-month period, is 4 percent *p.a.*

[5] This linear annualisation/deannualisation when interest is due more then once a year is the convention for bonds, notes and loans. It would have been more correct to annualise the effective rate, 2.5% per semester, via compound interest into $1.025^2 - 1 = 5.0625\%$ because we use compounding in the valuation too. But life is full of inconsistencies.

- The current two-month ZAR LIBOR is 3.5 percent *p.a.*

To value the (incoming) ZAR cash flows, note that there are eleven remaining interest payments at 3.5 percent each, the first of which is due two months from now. Discounted at 5/2 = 2.5 percent per half-year, the value is:[6]

$$\begin{aligned} \mathrm{PV}^{\text{fix}} &= 10m \times [1 + (0.035 - 0.025) \times a(2.5\%, 11)] \times 1.025^{4/6} \\ &= \text{ZAR}\, 11,133,193. \quad (7.9) \\ \mathrm{PV}^{\text{flo}} &= 10m \times \frac{1 + 1/2 \times 0.04}{1 + 2/12 \times 0.035} = \text{ZAR}\, 10,041,425. \quad (7.10) \end{aligned}$$

So the value of the fixed-rate leg exceeds the value of the floating-rate leg by

$$\text{ZAR}\, 11,133,193 - 10,041,425 = \text{ZAR}\, 1,691768. \quad (7.11)$$

This is what the company should receive for its swap contract.

7.3.2 Base Swaps

Under a base swap, the parties swap two streams of floating-rate interest payments where each stream is determined by a different base rate. For example, a LIBOR-based revolving loan can be swapped for a US T-bill-based revolving loan. The spread between these two money-market rates is called the TED spread (treasury-eurodollar spread). The TED spread is non-zero because T-bills and euro-CDs are not perfect substitutes in terms of political risk[7] and default risk. TED swaps can be used either to speculate on changes in the TED spread, or to hedge a swap book containing contracts with different base rates.

Example 7.9
The US office of a major bank has signed a fixed-for-floating swap based on the USD T-bill rate, while the London office has signed a floating-for-fixed swap based on USD LIBOR. This bank now has the USD T-bill rate as an income stream, and the USD LIBOR rate as an outgoing stream. To cover the TED-spread risk, it can swap its T-bill income stream for a LIBOR income stream using a base swap. The counterparty to this swap may be a speculator or simply another swap dealer who faces the opposite problem.

[6]The correction $1.025^{4/6}$ reflects the fact that the last coupon was four months ago; that is, we value the bond 4/6ths into the first coming coupon period, not at the beginning of that period.

[7]Dollar deposits in London cannot be blocked by the US Government, which is attractive to some parties. This is not a major issue anymore.

Figure 7.5: **Renault's 1981 circus swap**

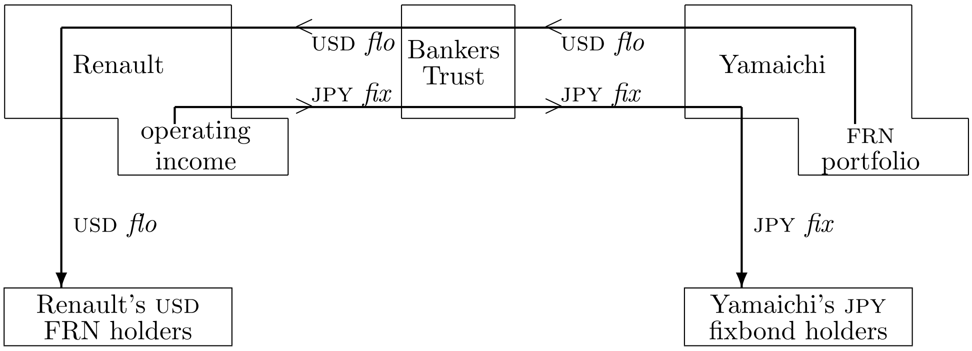

7.4 Cross-Currency Swaps

The cross-currency swap, or circus swap, is a currency swap combined with an interest rate swap (floating versus fixed rate), in the sense that the loans on which the service schedules are based differ by currency *and* type of interest payment. An early example is the Renault-Yamaichi swap already mentioned in Section 7.2.1. The historic background for the swap is as follows:

- Renault, a French car producer, wanted to get rid of its USD floating rate debt, and wanted to borrow fixed-rate JPY instead. The snag was that, because of Japanese regulations at the time, Renault was not permitted to borrow in the Japanese market.

- Simultaneously, Yamaichi Securities was being encouraged by Japan's Ministry of Finance to buy USD assets.[8] It could have bought, for instance, Renault's USD floating rate notes but was unwilling to take the exchange risk.

With the help of Bankers Trust, an investment bank, Renault convinced Yamaichi to borrow fixed-rate JPY and to buy floating-rate USD notes of similar rating and conditions as Renault's notes. As illustrated in Figure 7.5, Yamaichi was to hand over the USD service income from the USD investment to Renault, who would use the floating-rate USD interest stream to service its own floating-rate notes. As compensation, Renault undertook to service Yamaichi's equivalent fixed-rate yen loan, and pay a spread to both Bankers Trust and to Yamaichi.

The advantages of the swap to each party were:

- Renault was able to access the JPY capital market and get rid of its USD liability. (A more obvious solution would have been to borrow JPY and retire

[8]Japan wanted to show it was doing its bit to help finance the US deficit and also help "recycle the Petrodollars", a big issue after the second oil shock, early eighties.

the USD floating-rate notes with the proceeds. However, the first part of this transaction was not legally possible and the second part would have been expensive in terms of transaction costs or call premiums.)

- Yamaichi earned a commission. In addition, it now held USD assets (which was politically desirable) but these assets were fully hedged against exchange risk by the swap.
- Banker's Trust earned a commission on all of the payments that passed through its hands, plus a fee for arranging the deal.

The swap is also memorable because, even though it came quite soon after IBM-WB, it is already much more modern: it was not negotiated directly between two companies, but set up by Bankers Trust. Relatedly, there was no direct swap contract between Renault and Yamaichi, but two contracts (the *double swap*, as it was called then): Renault *v* Bankers Trust, and Bankers Trust *v* Yamaichi. This way, Bankers Trust took over the counterparty risk from Renault and Yamaichi, or, to be more precise, replaced the original counterparty risk by the risk that BT itself may get in trouble.[9]

7.5 CFO's Summary

The interest paid on any loan can be decomposed into the risk-free rate plus a spread that reflects the credit risk of the borrower. Swaps allow a company to borrow in the market where it can obtain the lowest spread, and then exchange the risk-free component of the loan's service payments for the risk-free component of another loan. This is useful if the other loan is preferred in terms of its currency of denomination or in terms of the way the periodic interest payments are determined (fixed or floating), as shown in Table 7.5. The use of risk-free rates within the swap is justified because the right of offset and the credit trigger eliminate virtually all risk from the swap—even though the company's ordinary loans remain risky. If desired, also the original loan's risk-spread payments can be swapped into the desired currency without altering their PV.

The difference between the spreads asked in different market segments usually reflects an information asymmetry—for instance, the firm's house bank often offers the best spreads because it is is less afraid of adverse selection—but may also reflect an interest subsidy. Another advantage is that the swap contract is a single contract, and is therefore likely to be cheaper than its replicating portfolio (borrowing in one market, converting the proceeds into another currency, and investing the resulting amount in another market). Other potential advantages include tax savings, or access to otherwise unavailable loans, or advantages of off-balance-sheet reporting.

[9] BT did get in trouble, and was taken over by Deutsche bank. Also Yamaichi sank ignominiously and was absorbed by Nomura, Japan's largest broker and investment bank.

Table 7.5: **Swaps: Overview**

preferred loan is ...		direct loan with best spread is ... fixed rate	floating rate
fixed rate	same currency	(do not swap)	do interest swap
	other currency	fix/fix currency swap	do circus swap
floating rate	same currency	do interest swap	(do not swap)
	other currency	do circus swap	flo/flo currency swap

Depending on the combination of the preferred type of loan and the cheapest available loan, one could use a fixed-for-fixed swap, a fixed-for-floating swap, or a circus swap. Each of these swaps can also be used to speculate on changes in interest rates or exchange rates. Likewise, base swaps are used to hedge against, or to speculate on, changes in the TED spread.

These four swaps are just the most common types; in fact, many more exotic swap-like contracts are offered. Thus, swaps have become increasingly popular with financial institutions and large corporations. All of these swaps are based on the principle of initial equivalence of the two legs of the contract. Thus, like forward contracts on exchange rates or currencies, the initial value of a swap is zero. To compute the market value of such a contract after inception, we just value each leg in light of the prevailing exchange rates and interest rates.

7.6 TekNotes

Technical Note 7.1 The value of a bond as a function of its yield to maturity

The valuation formula is derived as follows. Let the face value be 1, the coupon c per period, and the first coupon due exactly 1 period from now. (The yield is denoted by R, not r, because a compound per-period yield on a coupon bond should not be confused with an effective simple return on a zero-coupon bond.) We start with (almost) a definition; use the annuity formula; add and subtract 1; divide and multiply by R; and re-arrange:

$$\begin{aligned}
\text{PV} &= \sum_{t=1}^{n} \frac{c}{(1+R)^t} + \frac{1}{(1+R)^n} \\
&= c\,\frac{1-(1+R)^{-n}}{R} + (1+R)^{-n} \\
&= c\,\frac{1-(1+R)^{-n}}{R} + (1+R)^{-n} - 1 + 1 \\
&= c\,\frac{1-(1+R)^{-n}}{R} + R\,\frac{(1+R)^{-n}-1}{R} + 1 \\
&= (c-R)\,\frac{1-(1+R)^{-n}}{R} + 1 = (c-R)\,a(n,R) + 1. \qquad (7.12)
\end{aligned}$$

If the time to the next coupon is $1-f$ rather than unity, the PV rises by a factor $(1+R)^f$.

7.7 Test Your Understanding

7.7.1 Quiz Questions

1. How does a fixed-for-fixed currency swap differ from a spot contract combined with a forward contract in the opposite direction?

2. Describe some predecessors to the currency swap, and discuss the differences with the modern swap contract.

3. What are the reasons why swaps may be useful for companies who want to borrow?

4. How are swaps valued in general? How does one value the floating-rate leg (if any), and why?

7.7.2 Applications

1. The modern long-term currency swap can be viewed as:

   (a) a spot sale and a forward purchase.

   (b) a combination of forward contracts, each of them having zero initial market value.

   (c) a combination of forward contracts, each of them having, generally, a non-zero initial market value but with a zero initial market value for all of them taken together.

   (d) a spot transaction and a combination of forward contracts, each of them having, generally, a non-zero initial market value but with a zero initial market value for all of them taken together.

2. The swap rate for a long-term swap is:

   (a) the risk-free rate plus the spread usually paid by the borrower.

   (b) the risk-free rate plus a spread that depends on the security offered on the loan.

   (c) close to the risk-free rate, because the risk to the financial institution is very low.

   (d) the average difference between the spot rate and forward rates for each of the maturities.

3. The general effect of a swap is:

   (a) to replace the entire service payment schedule on a given loan by a new service payment schedule on an initially equivalent loan of another type (for instance, another currency, or another type of interest).

(b) to replace the risk-free component of the service payment schedule on a given loan by a risk-free component of the service payment schedule on an initially equivalent loan of another type (for instance, another currency, or another type of interest).

(c) to change the currency of a loan.

(d) to obtain a spot conversion at an attractive exchange rate.

4. You borrow USD 1m for six months, and you lend EUR 1.5m—an initially equivalent amount—for six months, at p.a. rates of 6 percent and 8 percent, respectively, with a right of offset. What is the equivalent spot and forward transaction?

5. Your firm has USD debt outstanding with a nominal value of USD 1m and a coupon of 9 percent, payable annually. The first interest payment is due three months from now, and there are five more interest payments afterwards.

   (a) If the yield at par on bonds with similar risk and time to maturity is 8 percent, what is the market value of this bond in USD? In Yen (at S_t = JPY/USD 100)?

   (b) Suppose that you want to exchange the service payments on this USD bond for the service payments of a 5.25-year JPY loan at the going yield, for this risk class, of 4 percent. What should be the terms of the JPY loan?

6. You borrow NOK 100m at 10 percent for seven years, and you swap the loan into NZD at a spot rate of NOK/NZD 4 and the seven-year swap rates of 7 percent (NZD) and 8 percent (NOK). What are the payments on the loan, on the swap, and on the combination of them? Is there a gain if you could have borrowed NZD at 9 percent?

7. Use the same data as in the previous exercise, except that you now swap the loan into floating rate (at LIBOR). What are the payments on the loan, on the swap, and on the combination of them? Is there a gain if you could have borrowed EUR at LIBOR + 1 percent?

8. You can borrow CAD at 8 percent, which is 2 percent above the swap rate, or at CAD LIBOR + 1 percent. If you want to borrow at a fixed-rate, what is the best way: direct, or synthetic (that is, using a floating-rate loan and a swap)?

9. You have an outstanding fixed-for-fixed NOK/NZD swap for NOK 100m, based on a historic spot rate of NOK/NZD 4 and initial seven-year swap rates of 7 percent (NZD) and 8 percent (NOK). The swap now has three years to go, and the current rates at NOK/NZD 4.5, 6 percent (NZD three years), and 5 percent (NOK three years). What is the market value of the swap contract?

10. Use the same data as in the previous exercise, except that now the NZD leg is a floating rate. The rate has just been reset. What is the market value of the swap?

Part III

Exchange Risk, Exposure, and Risk Management

About this Part

This Part is focused on the economics of exchange risk and hedging. To set the scene, we look into the question whether exchange rate changes are easy to understand and predict (Chapters 10 and 11). If so, there would not be much of a problem: all predicted changes would already be built into contracts, and there would be no bad surprises. Unfortunately, it turns out that exchange-rate movements are hard to predict; worse, they are even hard to understand and explain *ex post*.

We saw in Chapter 3 that real exchange rates can move a lot, and that this is important to firms. Coupled with the finding that most of the change comes from the nominal rate and that this part is hard to predict, it seems obvious that hedging is a good idea. We have to qualify, upon reflection: our conclusion in Chapter 12 is that hedging adds value if and only if it affects the company's real operations, not just its bank account.

Given that there are many channels through which the decision whether to hedge or not may affect operations, we conclude that hedging should often be relevant. The next question then is what size the forward hedge should be. What's the amount at stake? Chapter 13 reviews the various exposure concepts. Chapter 14 shows how to quantify the remaining unhedged risks as part of all market-related risks. We conclude with a review of ways to handle credit risk and transfer risk in international trade.

The minicase that follows brings up most of the issues.

Danish Weaving Machines

This is Copenhagen, in the late afternoon of 31/12/2005. Amidst the din of popping champagne corks, you (a trainee) and three regulars (Peter, Paul, and Mary) are still working hard. This very evening your firm, Danish Weaving Machines (DWM), has to submit its bid for an international tender for the delivery of a piece of fully automated weaving equipment. The customer, Taiwan Weaving Amalgamated (TWA) has invited bids in the currency of the bidder (DKK for your firm). There is only one serious competitor, France's *Équipements de Tissage* (ET). Due to a combination of luck and intelligence work involving, among the less unspeakable things, a rather expensive lunch in Paris, you know that ET has submitted a bid of EUR 2.8m. TWA will make up its mind on 1/4/2006, and will look at the price only—your and ET's equipment are embarrassingly similar. Production and delivery take a few weeks, payment would be by a Banker's Acceptance payable on sight and drawn on TWA's bank, First National of Taiwan, under a D/A documentary credit opened by First National via an L/C confirmed by your bank, Handelsbanken. The production cost would be DKK 18m. How should you set your price?

That looks easy to Peter: "For two months in a row, the EUR has been at the bottom of the ERM band (at DKK/EUR 7.5), and it cannot go any deeper. So we set our price at DKK 20.999m, somewhat below ET's price (EUR 2.8m $\times$ 7.5 = 21m). This leaves us a nice, sure profit of DKK 2.999m." Paul disagrees. "You must be out of your mind", he shouts. "It's decidedly in the cards that the DKK will revalue soon; and bankers tell me that, if and when there is a re-alignment, then by a time-honored ERM rule it will be by the cumulative inflation differential since the last re-alignment, that is, about 8%, to DKK/EUR 6.9. Just look at these forward exchange rates in the afternoon issue of *Børsen*:

spot	30d	60d	90d	180d	360d
7.50	7.30	7.25	7.20	7.15	7.10

If that's not half-predicting a lower EUR rate, I'll eat my hat." (Knowing Paul's hat quite well, the others look awed.) "IF there is a realignment, ET would win hands down. So we should set our price at DKK 19.319", Paul concludes, "somewhat below EUR 2.8m $\times$ 6.9 = 19.32, so that we win whatever happens. This still leaves us a profit of DKK 1.319m. This profit, unlike Peter's figure, is really safe; and 1.319m in the hand is better than 2.999m in the bush."

Mary is less than fully supportive: "Proverbs are for nitwits. What's "a" bird anyway? What about two humongous birds in the bush *v* a tiny, scruffy specimen in the hand? That is, how do you know that the PV of the risky but potentially high-payoff bid is lower than the PV of the riskfree one? You haven't even stated what the probabilities of a devaluation are. Nor have you explained how you set the discount rate as a function of the uncertainty, and how you defined the risk."

A thoughtful silence follows (apart from the continuous popping of champagne corks, elsewhere in the office). Fortunately for Peter, Paul and Mary, at this very moment the managing director comes in and takes them into his one-horse open sleigh to *Ensemble*, a (then) Michelin-star restaurant on Tordenskjoldsgade, thus leaving you, the trainee, with the problem. You have to fax TWA tonight, and the wrong decision would end your career at DWM.

Issues

1. What occult meanings & dark messages might be hidden in the cryptic phrase “payment would be by a Banker’s Acceptance payable on sight and drawn on TWA’s bank, First National of Taiwan, under a D/A documentary credit opened by First National via an L/C confirmed by your bank, Handelsbanken” ?

   Read Chapter 15 to find out. For current purposes, take this as meaning you get paid upon shipment of the machines. Using this interpretation, think of the following issues:

2. Suppose you want to reduce the uncertainty about the exchange-rate change. Is there any theory or type of information that would help reduce the uncertainty, or at least come up with a probability?

   Read Chapter 10 to find out.

3. What kind of exposure is there if we submit the high DKK price, and if we submit the low price: contractual, operating, or acccounting exposure?

4. Can one determine the size of the exposure, and, if so, what is the hedged value?

   Read Chapter 13 to find out about these two questions.

5. In choosing between the two alternatives, could any additional considerations play a role, or do we have enough information for the decision?

6. Suppose the optimal decision involves exchange risk. Does it make a difference whether you actually hedge, or is computing a hedged value as a tool in decision making all that matters?

7. Suppose that you read the Call for Tenders again, and, lo and behold, it says (in rather small print) that submitting a EUR bid is allowed.

   Your first reaction is that this does not help, since in the presence of a forward market any bid in EUR can be hedged into a DKK bid and vice versa. Then you realize that this hunch is clearly incorrect. Why? What was wrong with your initial hunch?

   Read Chapter 12 to find out about questions 5-7.

8. What would be the exposure if you submit a price in EUR, say 2.79985? What, if possible, is the hedged value? Would you use this option to quote a EUR price?

9. What does TWA gain by adding the EUR option to the Call for Tenders?

Chapter 12

(When) Should a Firm Hedge its Exchange Risk?

From Chapters 3 and especially 10-11 you may, I hope, remember that (*i*) deviations from purchasing power parity are sufficiently large and persistent so as to expose firms to real exchange rate risk, and that (*ii*) it is difficult to predict exchange rates. In earlier chapters we already described how forward or spot contracts can be used to reduce or even eliminate the effect of unexpected exchange rate changes on the firm's cash flows. We have not yet discussed the relevance of doing so. Thus, the central question that we address in this chapter is: Do firms add value when hedging their foreign exchange risk?

A key element in the discussion will be the zero-initial-value property of a forward contract: when the hedge is set up, its net asset value is zero. In light of this we can rephrase the question as follows: how can adding a zero-value contract increase the value of the firm? We will argue that hedging does add value if its effect is not just to add a gain or loss on the hedge but also to change something else in the firm, like decreasing the chances of financial distress. But there is a second question we need to address too, namely: if the hedge does add value, cannot the shareholders do the hedging if the firm does not? To this question we will answer that there are many good reasons why home-made hedging is not a perfect substitute for corporate hedging. The bottom line of this chapter is, however, not that hedging adds value (full stop): rather, we'd say that there are circumstances under which hedging helps, but these circumstances surely do not apply to all firms all the time.

In the first section of this chapter, we describe how/when hedging may achieve more than just adding a gain or loss on a forward or spot contract. In Section 2, we dismiss some bad reasons that lesser human beings occasionally advance in favor of, or against, hedging and some FAQs, starting with the issue whether companies can't simply leave hedging to the shareholders. Our conclusions are presented in Section 3.

12.1 The effect of corporate hedging may not just be "additive"

Hedging affects, quite possibly, the expected future cash flows of the firm, and it surely affects risk. How can we simultaneously take these two aspects into account? A finance person would immediately point out one excellent summary measure of the expected-value and risk effects of hedging: look at its net effect on present value. So in this chapter we adopt the Modigliani-Miller-style point of view that financial decisions should be rated on the basis of their impact on the company's market value.

In this light, then, one way to focus the discussion is to raise the zero-initial-value property of a forward contract: when the hedge is set up, its net asset value is zero. So we can rephrase the question as follows: how can adding a zero-value contract increase the value of the firm? One innocent answer would be that the zero-value property is a short-lived affair: almost immediately after being signed, the contract's value already changes. But the reply to this red herring is that one cannot even predict whether the value change will be for better or for worse. So, again: how can a contract add value if, roughly speaking, the chance that it acquires a negative value is fifty percent?[1]

The serious answer is that the zero-value property applies to the cash flows generated by a stand-alone forward contract: $\mathrm{PV}_t(\tilde{S}_T - F_{t,T}) = 0$. But the effect of hedging may very well be that the firm's other cash flows—anything that has to do with investing and producing and marketing and servicing debt etc—are affected by the hedge operation too. If (and only if) that is the case, hedging adds value—not because its own cash flow $\tilde{S}_T - F_{t,T}$ would have a positive net value in itself, we repeat, but because that cash flow has by assumption beneficial side effects on the firm's existing or future business. So the added value, if any, stems from a useful interaction between the hedge's cash flow and the other cash flows of the firm.[2]

This gets us to the real question: how can hedging interact with the firm's other cash flows? Below, we discuss (*i*) reduction of financial-distress costs, both *ex post*

---

[1] A variant of the above puzzle runs as follows. In an efficient market, the argument says, the gain from hedging or from any forward deal must have zero expected value, so that on average hedging does not help. This version of the puzzle is factually wrong: the forward rate is a biased predictor of the future spot rate, implying that $\mathrm{E}(\tilde{S}_T - F_{t,T}) \neq 0$. Also, the claim confines the effect of hedging to a purely additive one; but we already know that any value from hedging must stem not from $(\tilde{S}_T - F_{t,T})$, but rather from interactions with other cash flows. Lastly, the argument focuses on expectations, ignoring risks. One should look at PV instead.

[2] This echoes a argument that may be familiar from the Modigliani-Miller literature: one of the sufficient conditions for the irrelevance of the company's debt or pay-out policy is that the firm's "investments"—operations, really—are not affected. This assumption rules out interactions between the debt or pay-out decisions and the other cash flows. Many post-MM arguments question precisely this assumption—most prominently, Jensen's "free cash flow" theory and MM's tax theory.

Figure 12.1: **SS Silja Europe—from choppy waters to a safe haven**

Key On the left, the ship in its (red and white) Viking colors; right: the white and blue Silja version. Source: Wikipedia.

and *ex ante*; (*ii*) reduction of agency costs; (*iii*) lower expected taxes; and (*iv*) less noise in the profit figures.

12.1.1 Corporate Hedging Reduces Costs of Bankruptcy and Financial Distress

The most obvious route through which hedging can affect the firm’s prosperity is by decreasing its risk of financial distress. A firm is said to be in financial distress when its income is not sufficient to cover its fixed expenses, including financial obligations. The state of financial distress can lead to bankruptcy, which of course involves direct costs from reorganization or liquidation and the like. Large, uncovered exposures, combined with adverse exchange rate movements, may send a firm into insolvency and bankruptcy, or may at least contribute to such an outcome.

Example 12.1

In 1992, Rederi AB Slite, a Swedish shipping company that ran a ferry between Sweden and Finland for the Viking Lines, should have taken delivery of a very large ship. She had been ordered some years before from Meyer Werft in Papenburg, Germany. At the time of signing the purchase contract, Slite had decided not to hedge the DEM outflow because the SEK was tied to a basket in which the DEM had a large weight, and because the DEM was at a substantial forward premium relative to the SEK. However, by September 1992, Sweden had been forced to abandon the link between the SEK and the DEM, which had appreciated substantially against the SEK by the end of 1992. As a consequence of the appreciation of the DEM, Slite could no longer afford the ship (which was already painted in Viking’s red&white colors, see picture 1). So Meyer Werft kept it and soon managed to charter it to Viking Line’s rival, Silja Line, which repainted it (white, mostly, picture 2), named it Silja Europe and put it on the—you guessed it—Stockholm-Helsinki line.[3] A few

[3]Adding insult to injury, the world’s first floating McDonald’s restaurant was located onboard the Silja Europa from its maiden trip until 1996, Wikipedia tells us. But in 1996 the McDonalds

months later, Slite keeled over and went bankrupt.

Outright bankruptcy is costly because of the costs associated with liquidation. In the absence of these costs, Slite's shareholders would simply have lost control of the firm to the bondholders and banks, who would have carried on the business as before or who would have sold their ownership rights to others who, in turn, would have gone on running the firm as before. That is, in the absence of what Miller and Modigliani (MM) call bankruptcy costs, the event of insolvency would not have affected the value of the firm as a whole. In reality, of course, bankruptcy is costly; and the cost includes not only the fees paid to receivers, lawyers, assessors, and courts, but also the potential end of operations, loss of clientèle and reputation, and therefore liquidation at fire-sale prices rather than at going-concern value.

Example 12.2

In 2006, a company called Schefenacker that made mirrors for Mercedes and BMW and the like, got in trouble and had to go through a reorganisation. Bondholders lost over half of their stake, and Mr Schefenacker himself surrendered three quarter of his shares to debtors in lieu of repayment. The company even moved its HQs to the UK so as to be able to restructure under English law. Only the legal advisors were radiant, coming out EUR 40m the richer, which was almost 10 percent of the company's original debt.

In the same year, British Energy was an even greater bringer of Joy & Happiness to the legal crowd: with debts of GBP 1.2b (plus liabilities for taking care of spent nuclear fuel and decommissioning power stations) it paid GBP 121m for legal advice related to its restructuring. (*The Economist*, December 15 2007, p 67). Even in a relatively simple case like Northern Rock's, the English bank that skirted failure in the 2007/08 subprime mess, Deringer (a London law firm) made USD 20m from advice to the bank, Slaughter & May made USD 6m from advice to the Treasury, and Clifford Chance Linklaters made undisclosed amounts from working for third parties (*The Economist*, March 15 2008, p 78).

Costs of restructuring are soaring because financial structures are more complex now. Instead of e.g. three levels (senior, unsecured, and subordinated—once viewed as quite byzantine) we now see e.g. first-lien senior / second-lien senior / mezzanine / senior subordinated / junior subordinated. In each of these "classes" a majority has to approve the deal, giving each such class a veto right and, thus, endless possibilities of wrangling and blackmailing.

Example 12.3

Another car-parts maker, Meridian Automotive System of Michigan, took 20 months to organize. First-lien lenders had to yield part of their rightful takings to second-lien colleagues, which meant that seniority no longer meant seniority. In the case of

was closed down and replaced by Silja Line's own hamburger restaurant.

yet another car-parts firm, American Remanufacturers, second-lien lenders vetoed a proposal by first-lien lenders to refinance. Rather than paying them off, the first-lien group upped sticks and let the firm go bankrupt; neither class got anything, in the end. The two groups of lien holders "just shot each other", one lawyer said. Unusually (and disappointingly for the lawyers) the whole thing took just eleven days. (*Ibid.*)

But even before a firm actually goes bankrupt, the mere potential of future financial distress can already affect the operations and the value of the firm significantly. Thus, if hedging can reduce the volatility of the firm's cash flows, and hence the likelihood of the firm being in financial distress, hedging increases the firm's current value. Let us consider three specific links between the financial state of a firm and its real operations.

- **The Product Market and Reorganization Costs** Many firms sell products for which after-sales service is needed. The firms typically offer product warranties. A buyer's decision to purchase such products depends on his or her confidence in the firm's after-sales service. These firms sell more and must, therefore, be worth more the lower the probability of their going out of business. Hedging, by reducing the volatility of cash flows, decreases the probability of (coming uncomfortably close to) bankruptcy.

  Example 12.4
  When the US computer manufacturer Wang got into financial problems, one of Wang's customers noted that, "Before the really bad news, we were looking at Wang fairly seriously [but] their present financial condition means that I'd have a hard time convincing the vice president in charge of purchasing. At some point we'd have to ask 'How do we know that in three years you won't be in Chapter 11 [bankruptcy]'?" (Rawls and Smithson, 1990, p. 11).

- **The Labor Market and Wage Costs** Risk-averse employees are likely to demand higher wages if their future job prospects are very uncertain. In the event of bankruptcy, a forced change of job generally entails monetary and/or nonmonetary losses to employees. Thus, the employees want to protect themselves by requiring higher wages when working for a firm that is more likely to be in financial distress. If they do not get the risk premium, they quit—and especially the best ones, who can easily start elsewhere. This does *not* sound good.

- **The Goods Markets and Purchasing Costs** Risk-averse suppliers are similarly likely to demand cash upon delivery payment or, if they want to avoid even the risk of useless truck rides, cash before delivery. Trade credit would now be possible only in return for a big mark-up for default risk. Again, this is not the best a firm can overcome.

- **The Capital Market and Refinancing Costs** Loan covenants can trigger early repayment if the firm's income falls below a stated level, or credit lines can be canceled and outstanding credits called if there is a material deterioration in the firm's creditworthiness. To the extent that refinancing is difficult or costly when things do not look bright, it is wise for the firm to reduce income volatility by hedging. Costs associated with refinancing include not just an increased risk spread but also the hassle and distraction of transacting and negotiating, new restrictions on management, additional monitoring and reporting, and so on.

Financial distress costs are not the only link between hedging and the firm's operations. Following Jensen (1986), one could argue that also agency costs create such a link.

12.1.2 Hedging Reduces Agency Costs

Agency costs are the costs that arise from the conflicts of interest between shareholders, bondholders, and the managers of the firm. We will argue that these agency costs can affect the firm's wage bill, its choice of investment projects, and its borrowing costs. Hedging, by reducing the volatility of a firm's cash flows, can reduce the conflict of interests between different claimants to the firm's cash flows and can increase the firm's debt capacity and reduce its cost of capital.

One conflict is that between the managers of the firm and the shareholders. The source of the problem is that, through their wages and bonus plans, the wealth of the managers depends to a large extent on the performance of the firm. Since managers cannot sell forward part of their lifetime future wages in order to diversify, the only way that they can reduce the risk to their human wealth is to hedge the exposure by creating negatively correlated cash flows through positions in the foreign exchange, commodity futures, and interest futures markets. As argued below, "home-made" hedging (by shareholders or, here, by managers) is not a good substitute for corporate hedging because personal hedging is expensive and difficult. In addition, there is likely to be a maturity mismatch between the hedge and the exposed human wealth, which creates a ruin-risk problem similar to the one mentioned in connection with marking to market in futures markets (see Chapter 6). The reason for the mismatch is that affordable forward contracts are likely to have short maturities, while the wages that are exposed are realized in the longer run. The maturity mismatch between the short-term hedge and the long-term exposure becomes a problem when the value of human wealth goes up. Then, the short-term hedge triggers immediate cash outflows, while the benefits in terms of wages will not be realized until much later. That is, the personal hedge creates liquidity problems and, in the limit, may lead to personal insolvency.

For the above reasons, managers dislike hedging on a personal basis, and want the firm to hedge instead. If the firm does not hedge, managers can react in two ways. First, they are likely to insist on higher wages, as a premium for the extra risk they have to bear. Second, if the firm has investment opportunities that are very risky, the managers may refuse to undertake such projects even if they have a positive net present value. As the shareholders have imperfect information about the firm's investment opportunities or the management's diligence and motives, there is little they can do about these actions of the managers. Thus, the shareholders are better off if the firm hedges its exposures: this will automatically also hedge the managers' exposures, and thus make them look more kindly on the once-risky projects as well as their own pay checks.

Another example of agency costs is the conflict that arises between shareholders and bondholders in the choice of investment projects. This conflict arises because bondholders get (at most) a fixed return on their investment, while shareholders receive the cash left over after bondholders have been paid off. That is, the shareholders have a call option on the value of the firm, with the face value of the firm's debt as the option's strike price. The value of an option increases when the volatility of the underlying asset increases. (If this last bit is new to you, you did not properly read Chapter 8 on options.) Thus, in the case of a levered firm that is close to financial distress, shareholders may have an incentive to undertake very risky projects even if the project's net present value is negative. This *overinvestment problem* (Jensen and Meckling, 1976) arises if, due to increased uncertainty, the value of equity (the option on the future value of the firm as a whole) increases even though the current value of the firm as a whole goes down.

Example 12.5

A company has assets worth 60, currently invested risk-free, and debt with face value 50. For simplicity, assume risk neutrality and a zero risk-free rate. An investment opportunity arises where the investment would be 60, and the proceeds either 100 or 0 with equal probability. Therefore, the NPV is $(100+0)/2-60=-10$. But the shareholders might nevertheless be tempted by this plan because it would distribute more than enough wealth away from the bondholders and toward themselves:

decision	future outcome		resulting PV
don't invest	$V_1 = 60$:	bonds 50	bonds 50
		stocks 10	stocks 10
			total = 60
invest	Lucky: $V_1 = 100$	unlucky: $V_1 = 0$	
	bonds 50	bonds 0	bonds 25
	stocks 50	stocks 0	stocks 25
			total = 50

Obviously, if the shareholders are sufficiently ruthless to undertake this investment, the bondholders are worse off. Similarly, when a firm is close to bankruptcy, shareholders may have an incentive not to take on risk-reducing projects, even if these projects have a positive net present value. This *"debt overhang" underinvestment problem* (Myers, 1977) occurs if the current value of the firm goes up when the project is undertaken but the value of the option on the firm (the equity) goes down.

Bondholders, of course, recognize and anticipate these potential conflicts of interest and, therefore, adjust the terms of their loan appropriately. One way is to increase the interest charged on the loan. Another is to impose restrictions on management (the bond covenant) which requires costly monitoring by a trustee, slows down management, and may inadvertently even prevent good investments. Thus, if by hedging one reduces the variability of the firm's cash flows, one also reduces the potential for conflicts of interest associated with financial distress, and one thereby avoids the above extra costs of borrowing.

12.1.3 Hedging Reduces Expected Taxes

Hedging reduces expected cash flows if taxes are convex rather than linear functions of income. One example of a convex tax function is a progressive tax schedule, where the tax rate increases with income. In this case, smoothing the income stream will imply a lower average tax burden.

Example 12.6

Suppose that if income is USD 100, you pay USD 45 in taxes, while if income is USD 50, you pay only USD 20 in taxes. The expected tax when the earnings are USD 50 without risk then equals USD 20, while the expected tax is USD 22.5 when earnings are, with equal probability, either USD 100 or 0.

It may be argued that most countries' corporate tax rate schedules are, for all practical applications, flat. However, a more subtle type of convexity is created by the fact that, when profits are negative, taxes are usually not proportionally negative. In some countries, there are negative corporate taxes, but the amount refunded is limited to the taxes paid in the recent past. Such a rule is called *carry-back*: this year's losses are deducted from profits made in preceding years, implying that the taxes paid on these past profits are recuperated. Still, carry-back is limited to the profits made in only a few recent years, which means that negative taxes on losses are limited, too. In other countries, there is no carry-back at all. All one can do is deduct this year's losses from potential future profits (*carry-forward*), which at best postpones the negative tax on this year's losses.

Example 12.7

In Belgium, firms are not allowed to carry back losses. If a particular Belgian firm's profits are either EUR 35m or EUR 15m with equal probability, the expected profit is EUR 25m and the expected tax (at 30 percent) is EUR 7.5m. In contrast, if its profits are either EUR 100m or –EUR 50m with equal probability, the expected profit is still EUR 25m but now the expected tax is (EUR 100m $\times$ 0.3 + EUR 0)/2 = EUR 15m. It is true that the potential EUR 50m loss can be carried forward and deducted from subsequent profits, but these later tax savings are uncertain, and even if they were certain, there would still be the loss of time value.

Now consider a case where a firm is allowed to carry back its losses. Even in this case, excessive variability of income can affect the tax liability if the current losses are larger than the profits against which they can be set off. In the US, for instance, there is a three-year carry-back provision. Suppose that a particular firm's profits in the last three years amounted to USD 30m. If, for the next year, its profits are either USD 35m or USD 15m with equal probability, the expected profit is USD 25m and the expected tax (at 30 percent) is USD 7.5m. In contrast, if its profits are either USD 100m or –USD 50m with equal probability, the expected profit is still USD 25m but the potential loss now exceeds the profits made in the past three years. This means that in case of losses, the firm can recuperate the taxes paid on the USD 30m recent profits (that is, there is a negative tax of USD 30m × 30% = USD 9m), and the remaining USD 20m "unused" losses can be carried forward. Thus, the expected tax is [(USD 100m × 0.3) + (–USD 30m × 0.3)]/2 = USD 10.5m rather than USD 7.5m. It is true that the unused losses of USD 20m can be deducted from subsequent profits, but these later tax savings are uncertain, and even if they were certain, there would still be a loss of time value.

While the convex-taxes argument in favor of hedging is logically unassailable, you will probably agree that quantitatively this looks like a less important effect than the earlier ones (and especially financial distress), unless losses cannot be carried back nor forward—for instance because the company is not likely to survive anyway.

12.1.4 Hedging May Also Provide Better Information for Internal Decision Making

Multidivisional multinationals need to know the operational profitability of their divisions. By having each division hedge its cash flows, a multinational knows each division's operating profitability without the noise introduced by unexpected exchange rate changes. This may lead to better decision making and may, thus, lead to an increase in expected cash flows.

Of course, the same information can be obtained in different ways, and the alternatives may be cheaper. The firm could request that all divisions keep track of their contractual exposure at every moment, and could afterwards compute how profitable each division would have been if it had actually hedged. Nowadays, this just requires some programming. Another alternative, similar in spirit, is to shift all exchange risk towards a reinvoicing center. Under such an arrangement, a Canadian production unit, for instance, sells its output to a reinvoicing center on a CAD invoice, while a Portuguese marketing subsidiary buys these products from the center on a EUR invoice. In terms of information per subsidiary, this achieves the same objective as the subsidiary-by-subsidiary hedging policy. The corporation may then decide, on other grounds, whether or not the reinvoicing center should hedge the corporation's overall exposure.[4]

[4]If the reinvoicing center is instructed to hedge its exposure, this is likely to be cheaper than a policy where each subsidiary hedges its own exposures. First, the reinvoicing center can economize

Actual hedging entails a (small) cost, but as-if-hedged financial reporting is not costless either, and the corporation's operations may be too small to justify the fixed costs of a reinvoicing center. Thus, the bottom line is that the choice between actual hedging and as-if-hedged financial reporting or reinvoicing will depend on the circumstances.

12.1.5 Hedged Results May Better Show Management's Quality to Shareholders, and Pleases Wall Street

This argument is very close to the previous one: without exchange-rate-induced noise, one better sees the effect of management's decisions. The difference is that now the audience targeted by the clearer picture is the outside shareholder, not headquarters. Thus, the effect on value is more direct, and informal solutions like pro forma as-if-hedged financial statements would be confusing or not credible.

A related argument is that analysts and investment bankers like stable profits, as this makes prediction and valuation easier. A hedging policy would contribute to that.

We conclude with a review of some open issues.

12.2 FAQs about hedging

12.2.1 FAQ1: Why can't Firms leave Hedging to the Shareholders—Home-made Hedging?

Fans of the original MM article may remember that the options of home-made leveraging (or unleveraging) and home-made dividends play a big role in the argument. So we likewise ask the question, here, whether the firm cannot simply leave the hedging to the shareholders. There are many arguments saying that home-made hedging will not do, or not do as well as corporate hedging:

- The existence of financial-distress costs or agency costs is the most fundamental reason why "home-made" hedging is an imperfect substitute for corporate hedging. In reality, no individual shareholder can buy a contract that perfectly hedges against the costs of financial distress, like the loss of value when customers vote with their feet or employees flee. The problem, in short, is that the home-made hedge just produces the final cash flow $\tilde{S}_T - F_{t,T}$, and not the interactions with the firm's other business that provide the true advantage from hedging.

on hedging costs because it can "net" (clear) offsetting exposures. Second, there are likely to be benefits from specialization and scale economies. Third, the reinvoicing center is often located in a tax haven and simultaneously serves to reduce (or at least postpone) taxation on part of the group's profits.

- But even if hedging were purely additive, home-made hedging would not do as well as corporate hedging:
  - One reason is that, in the real world, shareholders have far less information than the managers about the firm's exposure. If shareholders have very imprecise knowledge of the firm's exposure, "home-made" hedging will be far less effective than corporate hedging.
  - Because of economies of scale, firms can obtain better terms for forward or money-market hedging than the individual shareholder. Thus, shareholders may value financial transactions undertaken for them by the firm.
  - Short-selling constraints can provide an additional reason why hedging is better undertaken by the firm rather than left to individual shareholders. In idealized markets, investors can easily borrow (or sell forward) any currency that they choose. However, in financial markets, personal borrowing in foreign currencies is not easy, and forward positions require substantial margin or else are discouraged by banks. It is true that going short is easy in futures markets; but the size of the futures contracts, however modest, may still be too large for shareholders with small positions in exposed equity. Moreover, for many currencies, there simply are no futures markets.

  Thus, corporations have better hedging opportunities than individual shareholders, which again means that "home-made" financial decisions are a poor substitute for corporate decisions.

12.2.2 FAQ2: Does Hedging make the Currency of Invoicing Irrelevant?

Does it matter whether prices are quoted in terms of the home currency or the foreign currency?

- The traditionalists state that someone must bear the exchange risk. Either you invoice in HC, in which case the foreign customer bears the exchange risk, or you invoice in FC, in which case you bear the exchange risk.
- The radical young turks believe that, with the existence of a forward market, there is no problem.

  Example 12.8

  Giovanni wants to buy his Carina GTI directly from Japan and calls Mr Toyota. We could envision two ways to set (and pay) the price:
  - In story 1, Mr Toyota ask JPY 2m 60 days. Giovanni agrees and immediately hedges at JPY/EUR 125 60 days. Thus, Giovanni's cost is locked in at 2m/125 = EUR 16,000 60 days.
  - Alternatively, Mr Toyota could ask EUR 16,000 60 days. If Giovanni agrees, Mr Toyota immediately hedges at JPY/EUR 125 60 days, and locks in an inflow of $16,000 \times 125 =$ JPY 2m 60 days.

So the currency of invoicing, in the young turks' view, merely shifts the hedging from seller to buyer, or vice versa. Finally, it does not matter which party hedges since, at a given point in time, each party can buy the foreign currency at the same rate.

While the above point of view is correct, you should realize that the example has two special features that are surely not always present. Notably, in the Toyota example the buyer and the seller are effectively able to hedge at the same moment and at the same rates. Conversely, the invoicing currency may matter as soon as (*i*) there is a time lag between the moment a price is offered by the exporter and the moment the customer decides to actually buy the goods, or (*ii*) the cost of hedging differs depending on who hedges. We illustrate these situations in the examples below. The first one focuses on the delay between the price offer and the customer's decision, the second one about differential costs:

Example 12.9

The currency of invoicing matters when you publish a list of prices that are valid for, say, six months. The problem here stems from the fact that there is a lag between the time that the FC prices are announced and the time the customer purchases an item. Since you do not know the timing and volume of future sales, you cannot hedge perfectly if you list prices denominated in foreign currency. Not hedging until you do know, on the other hand, may mean that by that time the rate has changed against you.

Example 12.10

The Argentina sales branch of a Brazilian stationery distributor instructs its customers to pay in BRL. Since the orders are frequent, and usually small, the Argentinian customers pay substantial implicit commissions whenever they purchase BRL. It would be cheaper if the exporter let them pay in Peso (ARS) and converted the total sales revenue into *Real* once a day or once a week.

In situations like this, one can still hedge approximately if sales are fairly steady and predictable. Many companies hedge all expected positions within a twelve-month horizon, and adjust their forward positions whenever sales forecasts are revised. However, in other cases, the time lag between the exporter's price offer and the importer's purchase decision may imply substantial sales uncertainty. In perfect markets, even this risk should be hedgeable at a low cost. In practice, the cost of hedging may very well depend on the currency in which prices are expressed.

Example 12.11

Here we consider an international tender, characterized by a time delay and a differential cost of hedging. Suppose that a Canadian hospital invites bids for a scanner.

1. **Buyer's currency** In an international tender, suppliers are usually invited to submit bids in the buyer's currency (CAD, in this case). A foreign contender's dilemma is whether or not to hedge, considering that:

- Forward hedging may leave the contender with an uncovered, risky forward position. Specifically, if the contract is not awarded to him, the bidding firm would then have to reverse: it would have to buy CAD spot—or forward, if the contract is reversed earlier—just to be able to deliver them, as stipulated in the forward contract. The rate at which such a time T purchase will be made is uncertain and can surely lead to losses.
- Not hedging at all means that, if the contender does make the winning bid, the CAD inflow is risky.

Thus, whether or not the contender hedges, there is a potential risky cash flow in CAD. It is true that banks offer conditional hedges, that is, contracts that become standard forward contracts (or standard options) when the potential supplier wins the tender but are void otherwise. However, these products are very much tailored to specific situations. The bank must assess and monitor the probability that a particular contender makes the winning bid, which makes such a contract expensive in terms of commissions. Thus, hedging is costly when bids are to be expressed in the customer's currency.

2. **Supplier's currency** The alternative is that the buyer invites bids in the suppliers' own currencies. Indeed, the buyer can easily wait until all bids have been submitted, then translate them into CAD $_T$—using the prevailing forward rates—and, at the very same moment she notifies the lucky winner, lock in the best price by means of a standard forward contract. In this way, all risk and all unnecessary bid-ask spreads in hedging disappear. To illustrate this, suppose that the Canadian hospital's procurement manager receives three bids in three different currencies, shown in column (a) below. She looks up the forward rates CAD$_T$/FC$_T$ shown in column (b), and extracts the following CAD$_T$ equivalent bid prices:

supplier	price	forward rate	CAD cost hedged
Oetker & Kölner, Bonn	EUR 120,000	CAD/EUR 1.65	CAD 198,000
Johnson Kleinwortsz, PA	USD 150,000	CAD/USD 1.35	CAD 195,000
Marcheix, Dubois & Fils, Québec	CAD 200,000		CAD 200,000

If price is the only consideration, she accepts the US offer, and immediately buys forward USD. Thus, when prices are to be submitted in the supplier's currency, a standard (and therefore cheap) forward hedge will suffice. ______

What this example shows, again, is that the currency of invoicing matters if the cost of hedging is not independent of the way prices are quoted. The Canadian hospital can use a cheap, standard contract if prices are submitted in the contending suppliers' home currencies. In contrast, with bids to be submitted in CAD, hedging is difficult and expensive for the bidders—because they are unsure about being awarded the contract. The solution in this case is to let the suppliers quote bids in their own currency. The general message to remember is that the option to hedge forward does not make the currency of invoicing irrelevant.

12.2.3 FAQ3: "My Accountant tells me that Hedging has cost me 2.17m. So how can you call this a Zero-cost Option?"

You accountant may have meant either of at least three things. First, she may have calculated that if you had not hedged, you would have raked in an extra 2.17m. Stated differently, the ex post sum of all the gains/losses ($\tilde{S}_T - F_{t,T}$) was –2.17m. This is sad indeed. But this is hindsight. All you can use, for decision-making purposes, is a PV criterion. And this brings us irrevocably back to the diagnosis that, in light of the zero-NPV property of ($\tilde{S}_T - F_{t,T}$), value stems from positive interactions, if any. The *ex post* value is just good or bad luck and is useless for decision making.

Alternatively, your accountant may have meant that the accounts show an *ex ante* cost of 2.17m. This concept is based on the not-infrequent (but misleading) practice of using spot rates to convert FC A/R's or A/P's into HC. If one then hedges, the actual cash flow differs from the book value, and the accountant hilariously calls this the cost of hedging. If, at the moment of booking the invoice, translation had been done at the forward rate, hedging would have entailed no accounting cost nor gain whatsoever.

Example 12.12

Recall our example in Chapter 5 of a Canadian firm that exports NZD 2.5m worth of goods. We were discussing an accounting issue: should we translate the A/R at the spot rate or at the forward? In that example we compared translation at the spot rate (0.90) and at the forward (0.88), and then looked at the outcome if the firm had not hedged. Now we assume the firm does hedge. The cost of goods sold being CAD 1.5m, profits then amount to 2.5m × 0.88 – 1.5m = 2.2m – 1.5m = 0.7m. But the operating profit depends on the initial valuation of the A/R, and the balance (if any) is called the cost/benefit of hedging:[5]

	using $S_t = 0.90$		using $F_t = 0.88$	
• **at** t**:**				
A/R	2,250		2,200	
COGS		1,500		1,500
operating income		750		700
• **at** T**:**				
bank	2,200		2,200	
hedging cost (D) or gain (C)	50		—	—
A/R		2,250		2,200

Which view is true? We know that hedging is free, in principle, so booking the

[5]In the old example, without hedging, there was a random capital gain. Here the effect is predictable, so accountants would call this a cost or benefit from hedging rather than a capital gain. Both are financial (i.e. non-operating) items.

forward premium as a cost or gain makes no sense. That accounting definition is a pure construct, based on the flawed practice of translating at the spot rate (see Chapter 4).

Of course, you could shrug off this accounting convention as irrelevant. There is, indeed, nothing wrong with writing weird things in books: the entire SF literature thrives on it. The only problem is that some people might actually believe this is a genuine cost in the same way *e.g.* the gasoil bill is a genuine cost. This risk arises especially among people that have no clue as to what accounting is about and simply believe a cost must be bad, otherwise it would be called a benefit.

In reality, there are in fact costs of hedging: there might be an upfront commission of a few Euros, and the bid-ask spread in the forward rate is always somewhat wider than in the spot. But these transaction costs have nothing to do with the forward premium, and they amount to a few basis points only.

12.2.4 FAQ4: "Doesn't Spot Hedging Affect the Interest Tax Shield, as Interest Rates are so Different Across Currencies?"

The last fallacy to be discussed is that hedging matters because it affects the interest tax shields. The issue is most often raised when the hedging alternative being considered is a money-market hedge rather than a forward transaction. Suppose, for instance, that a Russian company has accounts receivable denominated in Swiss Francs, and that the firm needs to borrow in order to finance its operations. CHF interest rates are much lower than RUB interest rates—say 6 percent as compared to 20 percent. If it borrows in RUB, the firm has a tax shield of 20 percent, and can reduce its taxes correspondingly. If it borrows in CHF, the loan also acts as a hedge, but the tax shield is a mere 6 percent. Thus, the argument concludes, the currency of borrowing affects the tax shields and, ultimately, the value of the firm.

As pointed out already in Chapter 4, the logical error in this argument is that it overlooks the fact that the taxes are affected not only by the interest paid, but (in the case of foreign currency borrowing) also by the capital gain or loss when the foreign currency depreciates or appreciates during the loan's life. Once this capital gain or loss is also taken into consideration, it is easily proven that, in PV terms, the currency of borrowing does not affect the current value of the firm even when there are taxes, as long as the tax on capital gains equals the tax on interest. Only when there is some form of tax discrimination may hedging affect the PV'ed tax shield.

Example 12.13

The UK used to have a rule that stated that exchange losses on long-term loans were deductible, but capital gains were tax-free. Given the risk-adjusted expectation that the AUD or NZD would depreciate relative to the GBP, a UK company had an incentive to borrow in, for instance, NZD or AUD. The expected capital gain would be tax-free, while the (then) high interest payments would be fully tax-deductible. Here there is a tax effect because taxes are discriminatory.

12.3 CFO's Summary

In the opening chapter of Part II we have argued that there are deviations from PPP. These deviations can be very large at any given point in time, and they also tend to persist over time. It typically takes three years before the distance between the actual spot rate and the PPP prediction is reduced by half. Moreover, it is difficult to predict exchange rates. All of this implies that firms that sell goods abroad, or import goods, or firms that compete with foreign firms or may have to compete with foreign producers in the future are exposed to real exchange rate risk. In this chapter, we have argued that it may be important that firms hedge this risk.

The Modigliani-Miller (1958) theorems state that financial policies, such as a firm's hedging strategy, cannot increase the value of a firm. However, this result is true only in perfect markets and if the firm's other cash flows are utterly unaffected by the financial decision at hand. Given the presence of convex tax schedules, costs of financial distress, and agency costs, hedging exchange risk can increase the value of a firm through its effect on future expected cash flows and the firm's borrowing costs. For a well-capitalized and profitable firm those considerations may carry little weight, and we do see many such firms happily ignoring exchange risk.

Not all companies are that lucky, though. For them, hedging adds value. But many comfortably rich companies have hedging policies too, often implemented by a reinvoicing center. Their view is that hedging may add little intrinsic value, but it is a low-cost option with some collateral attractions. For instance, managers like to reduce the risk of not meeting their numbers, Wall Street analysts appreciate predictibility, and HQ strategists prefer not to be distracted by items that have nothing to do with the division's own decisions. Also, strategists may argue that the decision *not* to hedge is not very different from a decision to speculate. There is nothing intrinsically wrong with speculation, but a firm's expertise is likely to be in its own business, not in speculating on foreign exchange. Thus, even thick-walleted companies often hedge their exposure.

12.4 Test Your Understanding

12.4.1 Quiz Questions

True-False Questions

1. In perfect markets, a manager's decision to hedge a firm's cash flows is irrelevant because there is no exchange rate risk.
2. In perfect markets, a manager's decision to hedge a firm's cash flows is irrelevant because the shareholders can hedge exchange risk themselves.
3. If a large firm keeps track of the exposure of each of its divisions, the firm has better information about each division, and is therefore better able to make decisions.
4. If a firm does not have a hedging policy, the managers may insist on higher wages to compensate them for the risk they bear because part of their lifetime future wealth is exposed to exchange rate risk.
5. If the firm does not have a hedging policy, the managers may refuse to undertake risky projects even when they have a positive net present value.
6. The risk-adjusted expected future tax savings from borrowing in your local currency always equals the present value of the expected tax savings from borrowing in a foreign currency.
7. The cost of hedging is roughly half of the difference between the forward premium and the spot exchange rate.
8. A reinvoicing center assumes the exchange rate risk of the various subsidiaries of a multinational corporation if it allows each subsidiary to purchase or sell in its "home" currency.

Valid-Invalid Questions

Determine which statements below are valid reasons for the manager of a firm to hedge exchange rate risk and which are not.

1. The manager should use hedging in order to minimize the volatility of the cash flows and therefore the probability of bankruptcy even though the expected return on the firm's stock will also be reduced.
2. Firms may benefit from economies of scale when hedging in forward or money markets, while individual shareholders may not.
3. The chance of financial distress is greater when a firm's cash flows are highly variable, and financial distress is costly in imperfect markets.

4. Shareholders do not have sufficient information about a firm's exposure.

5. Risk-averse employees demand a risk premium when the volatility of a firm's cash flows is high.

6. Short selling is often difficult or impossible for the individual shareholders.

7. Hedging a foreign currency inflow is beneficial when the forward rate is at a premium, because it is profitable and therefore desirable. In contrast, such hedging is not desirable when the forward rate is at a discount.

8. Since a forward contract always has a zero value, it never affects the value of the firm—but it is desirable because it reduces the variability of the cash flows.

9. Hedging reduces agency costs by reducing the variability of the firm's cash flows. Hedging means that the manager bears less personal income risk, making the manager more likely to accept risky projects with a positive net present value.

10. Hedging is desirable for firms that operate in a flat-tax-rate environment because income smoothing means that they can expect to pay less taxes.

11. Managers have an incentive to hedge in order to reduce the variability of the firm's cash flows because even though a firm may be able to carry forward losses, there is the loss of time value.

Multiple-Choice Questions

Choose the correct answer(s).

1. The Modigliani-Miller theorem, as applied to the firm's hedging decision, states that
   (a) in perfect markets and for given cash flows from operations, hedging is irrelevant because by making private transactions in the money and foreign exchange markets, the shareholders can eliminate the risk of the cash flows.
   (b) bankruptcy is not costly when capital markets are perfect.
   (c) a firm's value cannot be increased by changing the proportion of debt to equity used to finance the firm. Thus, the value of the tax shield from borrowing in home currency exactly equals the risk-adjusted expected tax shield from borrowing in foreign currency.
   (d) if the shareholders are equally able to reduce the risk from exchange rate exposure as the firm, then hedging will not add to the value of the firm.

(e) markets are perfect so hedging by the manager of the firm and the shareholders is irrelevant.

2. Hedging may reduce agency costs because
   (a) some of the uncertainty of a manager's lifetime income has been diversified away.
   (b) the shareholders will always prefer volatile projects while the debtholders will prefer nonvolatile ones.
   (c) risk-averse employees will demand a risk premium from a firm that is more likely to be in financial distress.
   (d) customers will think twice about purchasing goods from a company that may not be able to offer long-term customer service.
   (e) a reduction in the variability of the firm's cash flows may reduce the likelihood for conflicts between the debtholders and the shareholders.

3. Which of the following statements represent capital market imperfections?
   (a) Agency costs.
   (b) The difference between half of the bid-ask spread between the spot and forward markets.
   (c) The potential costs from renegotiating a loan that has gone into default.
   (d) The time value lost from having to carry forward losses into a future tax year.
   (e) Fees for liquidators, lawyers, and courts in the event of bankruptcy.

12.4.2 Applications

1. Using the following data, compute the cost of hedging for each forward contract in terms of implicit commission and in terms of the extra spread as a percent of the midpoint spot rate.

Maturity	Rates	Bid-ask	Hedging cost	Extra spread
Spot	49.858-49.898	0.040		
Fwd 30 days	49.909-49.965	0.056		
Fwd 60 days	49.972-50.043	0.071		
Fwd 90 days	50.061-50.157	0.096		
Fwd 180 days	50.156-50.292	0.136		

2. In the wake of the North American Free Trade Agreement, the firm All-American Exports, Inc. has begun exporting baseball caps and gloves to Mexico. Suppose that All-American is subject to a tax of 30 percent when it earns profits less than or equal to USD 10 million and 40 percent on the part of profits that exceeds USD 10 million. The table below shows the company's profits in USD under three exchange rate scenarios, when the firm has hedged its income and when it has left its income unhedged. The probability of each level of the exchange rate is also given.

	Hedged Profits	Unhedged profits	Probability
S_{hi}	15m	20m	25%
$S_{unchanged}$	10m	10m	50%
S_{lo}	5m	0	25%

(a) Compute the taxes that All-American must pay under each scenario.

(b) What are All-American's expected taxes when it hedges its income?

(c) What are All-American's expected taxes when it does not hedge its income?

3. In order to hedge its Mexican peso earnings, All-American is considering borrowing MXN 25 million, but is concerned about losing its USD interest tax shield. The exchange rate is USD/MXN 0.4, $r_{t,T} = 8\%$, and $r_t{}^*{}_{,T} = 6\%$. The tax rate is 35 percent.

   (a) What is All-American's tax shield from borrowing in USD?

   (b) What is All-American's tax shield from borrowing in MXN?

   (c) What is the risk-adjusted expected tax shield from borrowing in MXN?

4. Graham Cage, the mayor of Atlantic Beach, in the US, has received bids from three dredging companies for a beach renewal project. The work is carried out in three stages, with partial payment to be made at the completion of each stage. The current FC/USD spot rates are NZD/USD 1.6, DKK/USD 5.5, and CAD/USD 1.3. The effective USD returns that correspond to the completion of each stage are the following: $r_{0,1} = 6.00$ percent, $r_{0,2} = 6.25$ percent and $r_{0,3} = 6.50$ percent. The companies' bids are shown below. Each forward rate corresponds to the expected completion date of each stage.

Company	Stage 1	Stage 2	Stage 3
Auckland Dredging	NZD 1,700,000	NZD 1,800,000	NZD 1,900,000
forward rate NZD/USD	$F_{0,1} = 1.65$	$F_{0,2} = 1.70$	$F_{0,3} = 1.75$
Copenhagen Dredging	DKK 5,200,000	DKK 5,800,000	DKK 6,500,000
forward rate DKK/USD	$F_{0,1} = 5.50$	$F_{0,2} = 5.45$	$F_{0,3} = 5.35$
Vancouver Dredging	CAD 1,300,000	CAD 1,400,000	CAD 1,500,000
forward rate CAD/USD	$F_{0,1} = 1.35$	$F_{0,2} = 1.30$	$F_{0,3} = 1.25$

   (a) Which offer should Mayor Cage accept?

   (b) Was he wise to accept the bids in each company's own currency? Please explain.

Chapter 13

Measuring Exposure to Exchange Rates

We have established three important facts about the effect of exchange rate volatility on a firm's value. First, changes in the nominal exchange rate are not offset by corresponding changes in prices at home and abroad. That is, there are persistent and significant deviations from purchasing power parity, implying that there is real exchange rate risk (Chapter 3). Second, the forward rate is not successful in forecasting the exchange rate nor are other fundamental variables (Chapter 10). Third, given the market imperfections in the real world, hedging exchange rate risk can lead to an increase in the value of the firm (Chapter 12). We may conclude, therefore, that at least some firms may want to hedge their exposure to the exchange rate at least some of the time. The issue that is still unsettled is how much should be hedged. Specifically, one issue is whether hedging of contractual exposure, as discussed in Chapter 5, suffices: shouldn't we hedge all "expected" cashflows, whether contractual or not? And shouldn't we also think of the effect of exchange-rate changes on accounting values (as opposed to cash flows)?

In the first section of this chapter we distinguish between exchange-rate *risk* and *exposure to* the exchange rate. We next explain how one can classify the effects of exchange rate changes into two categories. First, exchange rate changes may have an impact on accounting values (known as *accounting exposure* or *translation exposure*). Second, the exchange rate may affect the firm's cash flows and market value (called *economic exposure*), either through its effect on existing contracts (labeled *contractual exposure* or *transaction exposure*) or through its impact on the future operating cash flows of the firm (known as *operating exposure*). Having already discussed the hedging of contractual exposures in Chapter 5, our discussion of this item here focuses on what it achieves, and where it stops, rather than on the mechanics (Section 2). The rest of the chapter then considers operating and translation exposure, in Sections 3 and 4, respectively.

13.1 The Concepts of Risk and Exposure: a brief survey

In general, we need to distinguish between the terms exchange risk and exchange exposure. (Some people use them interchangeably, which is not a good idea.)

- **Risk** We interpret *exchange risk* as synonymous with uncertainty about the future spot rate. Possible measures of exchange risk include the standard deviation or the variance of the future spot rate change.

- **Exposure** A firm is said to be *exposed* to exchange risk if its financial position is affected by unexpected exchange rate changes. A large exposure means that a given exchange rate change has a large impact on the firm. That is, by *exposure* we mean a numerical measure of how sensitive the financial position of a firm is to changes in the exchange rate.

This concept was already used in Chapter 5, where we generally defined exposure as a number that tells us by what multiple the HC value of an asset or cash flow changes when the exchange rate moves by ΔS, everything else being the same. We denoted this multiple by $B_{t,T}$:

$$B_{t,T} = \frac{\Delta\tilde{V}_T}{\Delta\tilde{S}_T}. \tag{13.1}$$

Note again the T subscripts to V and S: we have in mind values at T, so the delta's must mean that we compare two possible situations at the same (future) moment T, not two observations made at different moments in time. (We're so wont to interpret Δ as a change over time that explicit notation is in order, here.) Another way of saying this is that we have in mind a kind of *partial* derivative w.r.t. the exchange rate, holding constant other items (including time). We used this concept to price and hedge options (Chapter 9).

The above definition assumes that $\tilde{V}_T$ is an exact function of $\tilde{S}_T$. If the relation is known only up to noise or is otherwise imperfect—for instance because we willingly ignore non-linearities in the relation—a related concept of exposure crops up: the variance-minimizing hedge instead of the exact, perfect hedge. We looked at the variance-minimizing hedge already, notably in Chapter 6 on futures, and we'll use it again in this chapter. Recall that this hedge ratio is similar to the above exposure: a regression coefficient measures the sensitivity of $\tilde{V}_T$ to $\tilde{S}_T$, holding constant the regression residuals (which is "everything else", in a regression). So in that sense the general partial-derivative definition also covers the regression hedge-ratio measure of exposure.

We already showed that B has the dimension of a number of FC units. But what is meant by $\tilde{V}_T$? In the literature one typically lists three alternative possible specifications of what could be covered by this symbol:

- **Contractual exposure** In the case of contractual exposure, $\tilde{V}_T$ is defined as the HC value, at maturity, of a net contractual cash flow denominated in a

Figure 13.1: **Exposure Concepts: an Overview**

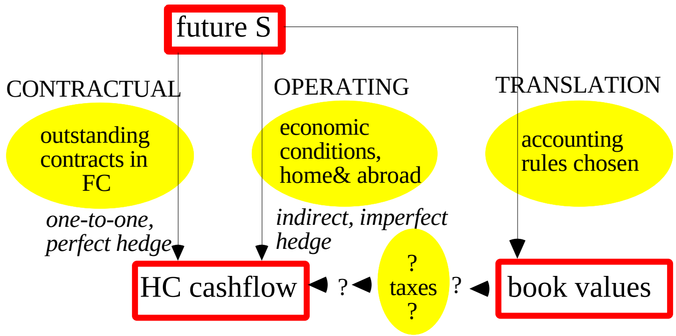

Key Contractual in- or outflows can be hedged peso-for-peso, if there really is no other risk. Operating exposures imply a noisy, convex relation between exchange rates and (HC-measured) cash flows, so the hedge is imperfect. Even getting a good idea of that relation is far from obvious: it requires a good understanding of the business and its environment. Translation exposure, lastly, primarily affects book values rather than cash flows—except indirectly if and when changed book values generate tax effects.

.

FC that matures on that date. It includes, per currency and per date, all A/R, A/P, deposits and loans denominated in a given FC, forward currency contracts, and contracts to buy or sell goods in future at known FC prices (Chapter 5). (Note that not all required information is found in the accounting system: commodity contracts where no delivery has been made yet, have not yet given rise to A/P or A/R, but they do generate contractual flows.) The exposure B then is the FC value, which is assumed to be risk-free.

- **Operating exposure** In the case of operating exposure one looks at the firm not as a portfolio of FC contracts signed in the past and generating cash in- or outflows in the future, but as a set of activities that require constant decisions by management, customers, and competitors. These future decisions depend, among other things, on future exchange rates, so that the cashflows are exposed in both FC and HC terms. (In the case of contractual exposure, in contrast, the FC amount is by assumption fixed, and only the HC value depends on the exchange rate.) In short, here $\tilde{V}_T$ is the cash flow from future operations rather than from past contracts, and the FC cash flow, C^*, is not a constant but depends on the future spot rate, $\tilde{S}_T$, and possibly other variables $\tilde{X}_T$: $\tilde{V}_T = C^*(\tilde{S}_T, \tilde{X}_T) \times \tilde{S}_T$.

- **Translation exposure** (or, less aptly, *accounting* exposure) arises when a multinational company has to consolidate its financial statements. As all subsidiaries' balance sheets and income statements are originally drawn up in the local currency, they must be translated first, and the result of this inevitably depends on the exchange rate at the reporting date.

Note that in the case of translation exposure we are talking about accounting values, that is, numbers written into books rather than cash flows that enter or leave bank accounts. To stress the difference, contractual and operating exposure are often referred to as *economic* exposures, as opposed to translation exposure.

We provide a more in-depth discussion of each of these in the rest of this chapter. We start with a discussion of contractual exposure.

13.2 Contractual-Exposure Hedging and its Limits

In Chapter 5 we saw already how one can close a contractual exposure, primarily by manipulating the financial items in the above list. We also saw how one can pool exposures for "similar" dates and hedge the aggregate net exposure, but too much grouping may create an interest-rate exposure problem. What we want to add now is a discussion of the limits and limitations of contractual-exposure hedging. First we consider the limitations: we show that hedging contractual exposure can achieve less than the uninitiated may have hoped. We then discuss the limits: how firm and certain should a cash flow be for it to be "contractual"; and what happens if we are less strict about this and include near-certain or even just "expected" cash flows?

13.2.1 What does Management of Contractual Exposure Achieve?

You may remember the example of Slite, the shipping line that keeled over when the devaluation of the FIM had made its new ship unaffordable. This could have been avoided by buying forward DEM. But this example is rather specific in that it involved a one-shot, and huge, exposure. The situation for a committed exporter or importer is different: there is a steady stream of in- or outflows, each of which is relatively small. The message to take home from this subsection is that even if such a firm continuously hedges all its contractual exposures, the impact of the exchange rate will be far from completely eliminated. There will still be exposure to the exchange rate from two sources: (*i*) exposure to variations in the forward rate, and (*ii*) "operating" exposure through the effect of the exchange rate on the volume of sales. We explain these issues below.

Consider an Italian firm, Viticola, which exports its fine wines to the US. Viticola can choose between at least two invoicing policies: (a) invoice in USD at (in the short run) constant US prices, and hedge each invoice in the forward market; or (b) invoice in EUR at (in the short run) constant home currency prices. In either case, Viticola has zero contractual exposure. Still, the exchange rate affects its profits:

- **Invoicing constant USD prices and hedging forward** Assume that the Italian firm extends three months credit to its US customers. If the firm hedges its contractual exposure systematically every time a new invoice is sent, its EUR cash flows ninety days later will be proportional to the ninety-day forward rate prevailing at the invoicing date. If, on the other hand, Viticola does not hedge its contractual exposure, its cash flows will be proportional to the spot rate prevailing when the invoice matures. In the long run, both series will have a similar variability, with the hedged version following the swings in the unhedged one with a three-month lag.

Example 13.1

Suppose that Viticola sets the price of a bottle of wine at USD 10. If Viticola does not hedge its transaction exposure, the revenue in EUR from US sales is random, and depends on the EUR/USD spot rate prevailing in three months time: USD $10 \times \tilde{S}_{\mathrm{t+3mo}}$. If, on the other hand, Viticola hedges each contract, the EUR cash flows from the sale of each bottle is USD 10 $\times F_{t,t+3mo}$. You should realize, however, that even though the forward rate for three months from now is known today, future forward rates are as uncertain as future spot rates. Thus, the revenue from future sales is an uncertain number, equal to USD $10 \times \tilde{F}_{T_i,T_i+3mo}$. Every decrease in the EUR/USD spot rate means a virtually identical decrease in the forward exchange rate, which then is reflected in lower revenue for Viticola three months later.

Thus, even perfect hedging of contractual exposure does not reduce the long-run variability of cash flows; it merely facilitates three-month budget projections.

- **Invoicing constant EUR prices** This means we let the exchange rate determine the USD price. From a contractual exposure point of view, Viticola is perfectly hedged since the contract is denominated in its home currency. Clearly, however, a policy of holding constant the domestic currency price may create huge swings in the USD price of the product and, therefore, may result in huge changes in the volume of Viticola's sales and profits, as illustrated below.

Example 13.2

Suppose that Viticola decides to set the price of each bottle of wine it sells at EUR 10. At the current spot rate of EUR/USD 1, this implies a price of USD 10, a price at which Viticola can sell 10,000 bottles in the US, and its total revenue from US sales is EUR 100,000. Assume that next month the USD depreciates to EUR/USD 0.95. Given that Viticola does not change its EUR price, the US price, translated at the new exchange rate, is now USD 10.53. At this new price, in the competitive wine market, Viticola can sell only 9,000 bottles. Thus, the export revenue of Viticola now declines to $9,000 \times 10 = 90,000$. True, the firm can now sell an extra 1000 bottles at home, but exports were the preferred solution (at the old rate, at least) and extra domestic sales probably require extra discounts too. Clearly, the total revenue of Viticola is exposed to the exchange rate.

The second policy, with its constant EUR prices, guarantees a stable profit per bottle sold but may cause big swings in volume. So the exposure is there, even if contractually there is none. The first policy, with its constant USD price, should guarantee fairly stable volumes, everything else being the same, but it leads to volatile profit margins. It is not obvious which of the two is the riskier, even after hedging. Hedging the expected USD revenue, if pricing is in USD, merely postpones the effects of exchange-rate changes on EUR revenue. In statistical jargon, hedging reduces the *conditional* variance of the 90-day cash flow to zero: conditional on what we know today (incl. the 90-day forward), there is no exchange-rate-related uncertainty about the 90-day cashflow. But unconditionally there is not much of a change in the variability. In still other words, Viticola's three-month budgets are less uncertain, but the uncertainty is merely pushed back 90 days. We still have no idea how the next three-month budget will look.

The alert reader may already have concluded that, in the long run, the pricing policy is actually more important than the invoicing decision. For instance, the exporter may invoice in EUR but adjust the EUR prices every month to compensate for changes in the exchange rate so as to keep the USD price roughly constant. In terms of contractual exposure, there is no risk (as invoicing is in EUR), but the variability of the profit margins remains. At the other extreme, the exporter may invoice in USD and hedge forward, but also adjust the USD price every month in order to maintain roughly constant EUR prices. Again, there is no contractual exposure, but the variability of the USD price and, hence, of the sales volumes remains. Whatever the policy, or whatever combination of policies a firm uses, future profits will remain exposed to exchange rate changes. Therefore, to hedge against changes in the exchange rate, one has to go beyond simply hedging contractual exposure.

13.2.2 How Certain are Certain Cashflows Anyway?

The other way to get to the same conclusion starts from the notion that the certainty seemingly implied by the word "contractual" is often illusory. There is always a non-zero probability of default on the counterpart's behalf, and occasionally the credit risk can be so big that one hesitates whether hedging is even a good idea.

Example 13.3

You signed a big export contract some time ago (time t_0), but now you hear that the company is in deep trouble. In fact, you estimate your chances of seeing the promised money to be about even. The deal is hedged and this forward sale has a current market value of $(F_{t_0,T} - F_{t,T})/(1 + r_{t,T})$. What to do now?:

- You could close out, "betting" on default by the customer. But if he survives and does pay, you have a open long spot position, the receivable.
- Alternatively, you could carry on, hoping for a happy end. The risk then is that there is default after all; and then you'll find yourself saddled with an open short forward position, this time the hedge.

Clearly, it is not obvious which alternative is more attractive: you are potentially damned if you do hedge and potentially damned if you don't. The only way to avoid dilemmas like this is to take out some form of credit insurance, which comes at a cost too.

While credit risk can be insured, other uncertainties about execution of a contract cannot. For instance, some contracts have built-in uncertainty, like cancellation clauses under certain conditions, or marking-to-market clauses if the exchange-rate change exceeds certain limits. In short, many contractual in- or outflows are not really certain.

On the other hand, some non-contractual positions are quite close to contracts, once one realizes that contracts offer no certainty anyway. What about a memorandum of understanding, or a letter of intent? What about a verbal deal—legally a contract as there is consensus, but hard to prove and, therefore, hard to enforce? What about near-certainty about future sales contracts based on experience from the past? Many committed exporters or importers would be tempted to go beyond pure contractual positions, and hedge also near-certain forex revenue, hoping to thus postpone the impact of exchange-rate changes beyond the credit period.[1]

13.2.3 Hedging "Likely" Cashflows: what's new?

One should realize that the hedging of "likely" cash flows has two implications. First, noise creeps in, stemming from other variables than future exchange rates. Second, abstracting from noise, the relation between the HC cash flow and the exchange rate is likely to be convex. That is, we go from an exact linear relation (like $\tilde{V}_T = B_{t,T}\tilde{S}_T$) to a noisy and non-linear one: $\mathrm{E}(\tilde{V}_T|S_T) = f_{t,T}(\tilde{S}_T)$. How come?

The *noise* comes from the fact that the final decision still is to be taken by the customer (or the exporter), and this decision will inevitably depend on other variables than the exchange rate. A car exporter's foreign sales, for instance, will depend on other producers' prices and promotions, on interest rates for personal loans, the level of consumer confidence, etc. The *convexity*, on the other hand, stems from optimal reaction to exchange-rate changes. The exporter does have the option to sell at a constant FC price, in which case the translated revenue would rise or fall proportionally with the exchange rate, everything else being the same. But this passive policy will be abandoned if the exchange-rate change is sufficiently big and if reaction does improve the situation. Thus, in 1974 VW might have been exporting its beetles to the US at USD 2,000 apiece, but with a falling dollar and shrinking profit margins they would surely increase the USD price if that beats the passive policy. (This should probably have come with further changes in the marketing mix.) Even

[1] Recall that pure A/P hedging just postpones the impact of shifts by the credit period, like three months in the Viticola example.

Figure 13.2: **How Convexity Arises in Operating Exposure**

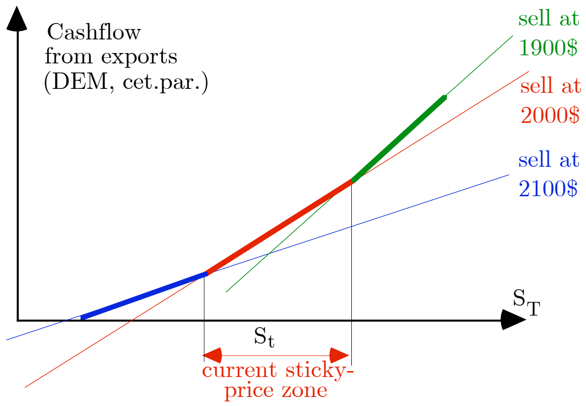

Key The lines show an exporter's HC cash flows for a given FC sales price, everything else being constant. The optimal price depends on the exchange rate. A policy of always choosing the best price leads to a convex relation between S and the expected cash flow.

abandoning exports would be an option: zero cash flows are better than negative ones. In the case of a rising dollar, similarly, VW might have considered lowering its USD price below 2,000, giving up some profit margin in exchange for more market share. Again, this will be adopted if it beats the passive policy. The final picture is one of a piecewise-linear, convex relation (Figure 13.2: passive sticky-prices policies for exchange rates close to the current level, but switching to new and better policies if the change has become sufficiently big.

In fact, in the above paragraph we have actually wandered from the realm of contractual exposure into that of short-term operating exposure. Before we proceed with this, let's point out one major implication of the fact that the effect of exchange-rate changes now is of a general non-linear form, $\mathrm{E}(\tilde{V}_T|S_T) = f_{t,T}(\tilde{S}_T)$, rather than a contractual-exposure type relation $\tilde{V}_T = B_{t,T}\tilde{S}_T$. The implication is that exposure is no longer some number of FC units that can be found in the balance sheet, or a FC cash flow as stated in a pro forma P&L. Rather, exposure has to be computed—notably from a comparison of two or more possible outcomes for the firm at time T, one outcome per possible exchange rate. As a colleague put it, "the idea is completely foreign to accounting-tied CFOs". Here's your chance to get ahead.

13.3 Measuring and Hedging of Operating Exposure

While contractual exposure focuses on the effect of the exchange rate on future cash flows whose value in foreign currency terms is contractually fixed in the past, operating exposure analyzes the impact of future exchange rates on *noncontractual* future cash flows. These FC cash flows that are likely to be random even in terms of the foreign currency, partly as a result of other factors than exchange rates and partly because of the exporter's endogenous response to the exchange-rate change. Thus, the complicating factors relative to contractual exposure are that the relation between the HC cash flow $\tilde{V}_T$ and the exchange rate $\tilde{S}_T$ has become noisy and non-linear. Worse, the relation has become hard to identify, as it depends on the economic environment that the firm competes in, and on how the firm reacts to changes in the exchange rate, given its competitive environment.

13.3.1 Operating Exposure Comes in all Shapes & Sizes

There are at least two misconceptions about the source of operating exposure. The first misconception, already discarded in the previous section, is that if a firm denominates all of its sales and purchases in terms of its own currency, it faces no exposure to the exchange rate. We know better, now. The second misconception is that only those firms that have foreign operations are exposed to the exchange rate; that is, only those firms that buy or sell goods abroad or use imported inputs are exposed to the exchange rate, while firms that have only domestic operations are not exposed to the exchange rate. This is usually wrong too. For instance, an exchange-rate change can turn a potential foreign exporter into an active competitor:

Example 13.4

Consider a firm located in the US. Assume that the firm's production is based in the US, and that the firm uses only inputs that are produced in the US and that the firm's entire sales are in the US. The naive view would suggest that this firm's operations are not exposed to the exchange rate. This view is false if the firm faces competition from abroad. Every time the USD appreciates, the foreign competitors gain; they can lower their USD prices and still obtain the same amount of their own home currency. US firms that faced this type of situation include Caterpillar, Kodak, General Motors, and Chrysler. In the early 1980s, when the USD appreciated against the JPY, all of these firms lost market share to their Japanese competitors, Komatsu, Fuji, Honda, and Toyota respectively. This erosion of market share led to large decreases in profits for the US firms.

The second way an apparently non-international player may be affected by exchange rates is indirectly, at a remove: the firm may buy from local firms that, in turn, do import; or it may sell to local firms that, in turn, do export. Or, even more indirectly: in an economy with a large open sector, the general level of economic activity may depend on the state of health of the export and the import-substituting industries.

Example 13.5

A UK firm has set up a subsidiary in our favorite country, Freedonia. Assume, for simplicity, that the subsidiary's cash flow, in terms of the Freedonian crown (FDK), can take on one of two (equally probable) values, FDK 150 or FDK 100, depending on whether the Freedonian economy is booming or in a recession. Let there also be two equally probable time-T spot rates, GBP/FDK 1.2 and 0.8. Thus, measured in terms of the home currency, the GBP, there are four possible outcomes for the future cash flows, as shown in Table 13.1. In each cell, we also show the joint probability of that particular combination of outcomes for the exchange rate and the state of the economy. When the FDK is expensive, a recession is more probable than a boom because an expensive currency means that Freedonia is not very competitive. The inverse happens when the crown is trading at a low level. Thus, we assume that the probability of the exchange rate being high and the economy booming is fairly low: 0.15 not 0.25,[2] and likewise for the unexpected combination of a cheap Krone and a slumping economy. The more expected outcomes get probabilities 0.35.

One step towards quantifying the impact of the exchange rate is to first compute the conditional expected cash flow for each level of the exchange rate—each row, in the table. These numbers are added in the right-most column of the table and amount to 138 when the rate is high, and 108 when the rate is low. Thus, the expected impact of the exchange rate change is 30 (million) pounds.

In the example, there is more risk than just the uncertainty about the exchange rate (with its differential impact of 30): here, there is no one-to-one relation between the state of the economy and the level of the exchange rate, so the the firm's cash flow is not yet fully certain once you observe (or hedge) the exchange rate. In regression parlance, this would be called a residual uncertainty.

The example also illustrates how the relation between the HC cash flow (or the FC cash flow) and the exchange rate can be noisy. Below, we give a simple example where a convexity arises from the exporter's optimal reaction.

Example 13.6

A French niche producer of bottled mineral water can export its output to the US, where it sells at USD 1.25 per bottle (the market price minus the shipment costs etc). But it can also sell at home, at EUR 1.00. Obviously, for $\tilde{S}_T < 0.80$, they better sell at home, while for higher rates the wiser solution is to export:

$$\tilde{V}_T = \begin{cases} 1.00 & , \text{ if } \tilde{S}_T \leq 0.80 \\ \\ 1.25 \times \tilde{S}_T & , \text{ if } \tilde{S}_T > 0.80 \end{cases} \tag{13.2}$$

So the function is a piecewise linear one. (Figure 13.3).

[2] If the health of the Freedonian economy had been independent of the level of the spot rate, the probability of each cell would be $0.5 \times 0.5 = 0.25$.

Table 13.1: **Joint distribution of $\tilde{S}_T$ and $\tilde{CF}_T$ for the Freedonian Subsidiary**

	boom: CF*= 150	bust : CF*= 100	$\mathrm{E}(\tilde{V}_T \mid S_T)$
$S_T = 1.2$	$150 \times 1.2 = 180$ $p = 0.15$	$100 \times 1.2 = 120$ $p = 0.35$	$\frac{180\times0.15+120\times0.35}{0.15+0.35} =$ GBP 138 $p = 0.50$
$S_T = 0.8$	$150 \times 0.8 = 120$ $p = 0.35$	$100 \times 0.8 = 80$ $p = 0.15$	$\frac{120\times0.35+80\times0.15}{0.35+0.15} =$ GBP 108 $p = 0.50$
	$p = 0.50$	$p = 0.50$	

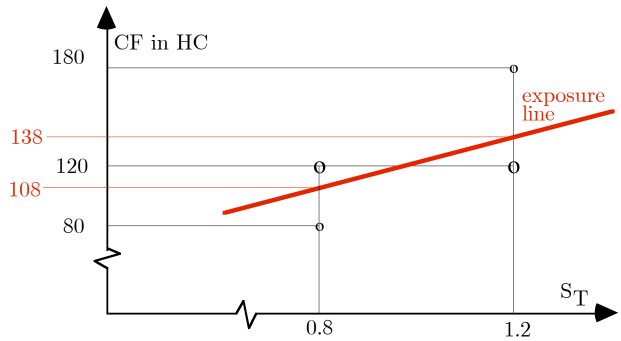

The above examples are all about short-term exposures. By short term we mean, like in micro-economics, that the investments (P&E) are given; no major expansion or downsizing or relocation is being considered. Recall the example where VW was revising its marketing and pricing policies in light of the DEM/USD exchange rate. These were short-term reactions. But VW's reaction did become "long-term" when it considered moving its production abroad. In the late 70s, it effectively built factories in Brazil, Mexico and the US.

Figure 13.3: ***Bourbonnaise des Eaux*'s Option to Export Mineral Water**

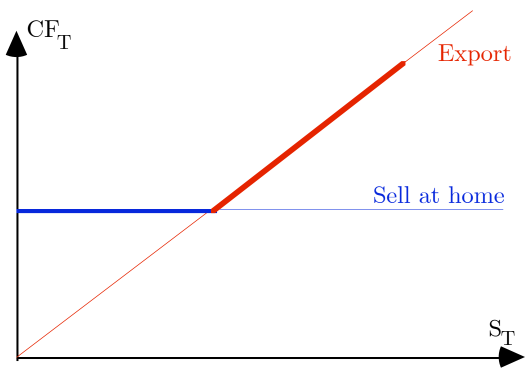

Thus, operating exposure comes in all kinds of shapes & sizes. How, then, can one still hedge it? What is the measure of exposure? This depends on the type of hedge instrument one has in mind. When hedging is done with a linear tool like a spot or forward position, we have to approximate the (noisy, non-linear—remember?) relation by a linear one, using regression. If a non-linear hedge is used, for instance a portfolio of options, things are different. We begin with linear hedges.

13.3.2 The Minimum-Variance Approach to Measuring and Hedging Operating Exposure

Note from the definition of operating exposure given in Section 3 that exposure tells us by how much the cash flows of the firm change, for a unit change in the exchange rate. Adler and Dumas (1983) suggest the use of simulations to compute the economic exposure. The simulation requires that we come up with a number of possible future values for the spot exchange rate and compute the value, in home currency, of the cash flows for each possible future exchange-rate value. The exposure of the firm to the exchange rate can then be computed by decomposing the HC value of the asset or cash flow, $\tilde{V}_{T,s}$ in scenario $s = 1, ...n$, into a part linearly related to the spot rate in that scenario and a part uncorrelated with the spot rate—a technique commonly called linear regression:

$$\begin{aligned}\tilde{V}_{T,s} &= A_{t,T} + B_{t,T}\tilde{S}_{T,s} + \tilde{\epsilon}_{t,T,s}, \\ &= \underbrace{A_{t,T} + \tilde{\epsilon}_{t,T,s}}_{\substack{\text{uncorrelated} \\ \text{with } \tilde{S}}} + \underbrace{B_{t,T}\tilde{S}_{T,s}}_{\substack{\text{exactly linear} \\ \text{in } \tilde{S}}} \quad .\end{aligned} \tag{13.3}$$

If $\tilde{V}$ were truly linear in $\tilde{S}$, we could have used the familiar conditional-expectation equation, $\mathrm{E}(\tilde{V}_{T,s}|S_{T,s}) = A_{t,T} + B_{t,T}S_T$, but that is usually not appropriate: the above is just a linear approximation or linear decomposition or linear projection of something that is really non-linear. But we need the linear approximation rather than the true relation because our hedge instrument is linear anyway. We start with a number of examples where the situation is so simple that the regression can be done naked-eye, almost. In the first illustration there isn't even any noise ($\tilde{\epsilon}$):

A Problem with just two Possible Exchange Rates, no Noise

Example 13.7

Belgium's Android MetaProducts NV/SA wishes to hedge its exposure to the exchange rate stemming from its ownership of a marketing affiliate located in UK. This is 1992, and the GBP has just formally joined the ERM after maintaining a constant rate for two years. Still there is risk: what worries Android is that, in the past few years, inflation has been substantially higher in the UK than on the continent, raising the question whether the current exchange rate, BEF/GBP 60, is sustainable. After discussion with its bankers, Android ends up with two possible outcomes:

Figure 13.4: **Android and the Pound**

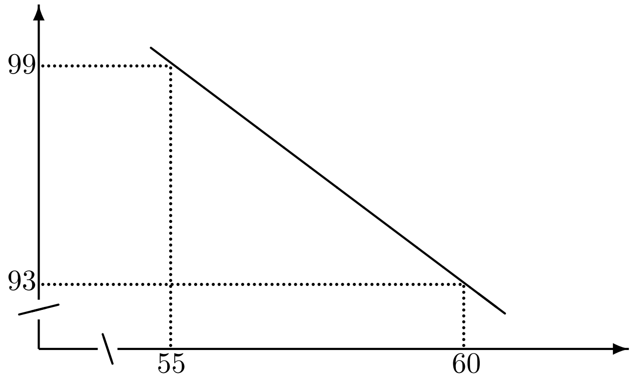

- The UK government may switch to a strongly deflationary policy and stabilize the exchange rate at 60. Such a deflationary policy is expected to depress sales and would decrease the net cash flow of the marketing affiliate to GBP 1.55m.
- Alternatively, the UK government may let the GBP depreciate and follow a moderately deflationary policy. In this case, the exchange rate would be BEF GBP 55, and management expects a cash flow of GBP 1.8m.

How can we hedge this? Obviously, as we have an asset denominated in pounds, the exposure seems bound to be positive—but should we hedge the lower amount, or the higher one, or something in between? The message below will be that the above "obvious" diagnosis is totally off the mark: the exposure is nowhere near the 1.55-1.80m range. In fact, it is massively negative. We see this by computing the two possible HC values:

$$\begin{aligned}\text{(no devaluation:)} \quad V_T &= 1.55m \times 60 = \text{BEF}\,93m,\\ \text{(devaluation:)} \quad V_T &= 1.80m \times 55 = \text{BEF}\,99m.\end{aligned}$$

This tells us that we win if the pound loses value, which means the exposure is negative. Figure 13.4 illustrates this. It is quite easy to compute the slope of the line connecting the two possible outcome points:

$$\text{slope} = \frac{93m - 99m}{60 - 55} = \frac{-6m}{5} = -\text{GBP}1.2m. \tag{13.4}$$

This slope is, of course, none other than our exposure, B: if there are just two possible points, the regression is the line through those two points. We now show that if Android takes a position in the forward market with the opposite sign—minus minus 1.2m, that is, buying forward 1.2m—it is hedged. Suppose that the forward rate is 58. The outcomes are analyzed as follows:

case	raw cash flow	outcome of hedge	hedged cash flow
S=60	93m	$1.2m \times (60 - 58) = +2.4m$	$93m + 2.4m = \text{BEF}95.4m$
S=55	99m	$1.2m \times (55 - 58) = -3.6m$	$99m - 3.6m = \text{BEF}95.4m$

DoItYourself problem 13.1

Verify that if the forward rate had been different, the level of the hedged cash flow would be affected but not the fact that the investment is hedged. For instance, with a forward rate of 57 instead of 58 the hedged asset would have been 1.2m higher, at 96.6m. Show it.

Remember two things from the example. First, exposure is computed from a comparison of alternative future outcomes, not from one single number found in a balance sheet or a pro forma cash flow statement for next year. Second, the size and (here) even the sign of exposure can be very different from what gut feeling would suggest. Here, an accounting-tied CFO would have taken for granted that exposure is positive: we talk about a GBP asset, don't we? Wrong; the position behaves like a 1.2m liability.

DoItYourself problem 13.2

We just showed that exposure defined as a slope of the line linking the two points does work: in this (overly simple) example, all risk is gone. Show that if you would have followed your intuition and had hedged (sold forward) GBP 1.55m, or GBP 1.8m, or in fact any positive number, the uncertainty after such "hedging" would have been higher than before.

A Problem with two Possible Exchange Rates and Noise

Let us generalize. The fact that the regression hedge always succeeds in taking away all exchange-related risk can be proven in just two lines:

$$\begin{aligned}\tilde{V}_T^{\text{hedged}} &= \underbrace{A_{t,T} + B_{t,T}\tilde{S}_T + \tilde{\epsilon}_{t,T}}_{\text{value unhedged}} + \underbrace{[-B_{t,T}]}_{\text{size of hedge}} \underbrace{\left[\tilde{S}_T - F_{t,T}\right]}_{\text{expiry value per GBP}}, \\ &= \underbrace{A_{t,T} + \tilde{\epsilon}_{t,T}}_{\text{uncorrelated with } \tilde{S}} + \underbrace{B_{t,T}F_{t,T}}_{\text{risk-free}} . \qquad (13.5)\end{aligned}$$

Thus, what regression-based hedging generally achieves is eliminating all uncertainty that is linearly related with the exchange rate: $B_{t,T}F_{t,T}$ has taken the place of $B_{t,T}\tilde{S}_T$. The uncertainty that is *not* correlated with the exchange rate, in contrast, cannot be picked up by the forward contract, so it remains there. It can be shown that this regression hedge ratio is also the one that reduces the variance of the remaining risk to the lowest possible level. This is why this section is called Minimum-Variance hedging and why ordinary regression is called Least Squares (=minimal residual variance).

In the Android example there was assumed to be no residual risk, which is hard to believe. Our earlier Freedonia example, in contrast, does have this feature: the state

Table 13.2: **Joint distribution of $\tilde{S}_T$ and $\tilde{CF}_T$ for the Freedonian Subsidiary**

	unhedged cash flows boom: CF*= 150	bust : CF*= 100	$E(\tilde{V}_T \mid S_T)$
$S_T = 1.2$	$150 \times 1.2 = 180$ $p = 0.15$	$100 \times 1.2 = 120$ $p = 0.35$	$\frac{180\times0.15+120\times0.35}{0.15+0.35}$ = GBP138 $p = 0.50$
$S_T = 0.8$	$150 \times 0.8 = 120$ $p = 0.35$	$100 \times 0.8 = 80$ $p = 0.15$	$\frac{120\times0.35+80\times0.15}{0.35+0.15}$ = GBP108 $p = 0.50$
	hedged cash flows		
$S_T = 1.2$	$180 - 18 = 162$	$120 - 18 = 102$	$\frac{162\times0.15+102\times0.35}{0.15+0.35}$ = GBP120
$S_T = 0.8$	$120 + 12 = 132$	$80 + 12 = 92$	$\frac{132\times0.35+92\times0.15}{0.35+0.15}$ = GBP120

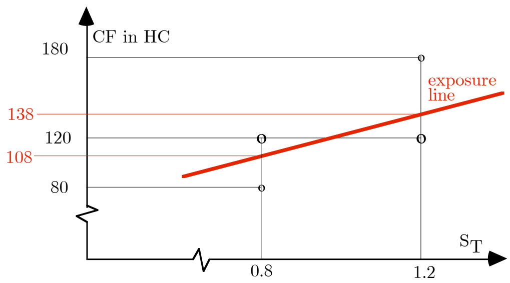

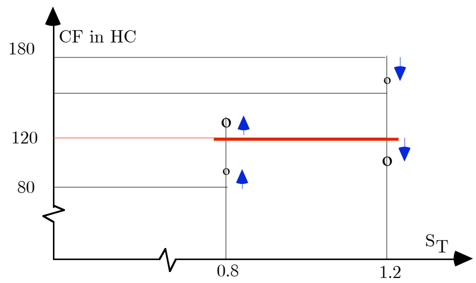

of the economy (and, therefore, the cash flow) is not fully know once the exchange rate is observed, so there is only an imperfect correlation between the HC cash flow and the exchange rate. Table 13.2 repeats the Freedonia data and then shows the hedged cash flows. To find the hedged cash flows we of course need the exposure. In the case with just two possible values of $\tilde{S}_T$, the regression line runs through the points representing the conditional expectations. We identified these expectations as 138 when $S_T = 1.20$ and 108 when $S_T = 0.80$. So the exposure now equals

$$B_{t,T} = \frac{138 - 108}{1.20 - 0.80} = \frac{30}{0.4} = \text{FDK}75. \tag{13.6}$$

Note, in passing, that even though the cash flow, in FC, is either 150 or 100, the exposure is not even in the range [100, 150]: it equals 75. The only way to come with a meaningful exposure number again is to compare the two scenarios; neither scenario in itself gives you a reliable answer, nor does any accounting number. Let's show that our FC 75 does make sense. Assuming the forward rate is 0.96, the pay-offs from the hedges would be

$$\begin{aligned}
\text{when } S_T = 1.20\text{:}\quad & -B_{t,T}(S_T - F_{t,T}) = -75 \times (1.20 - 0.96) &= -18,\\
\text{when } S_T = 0.80\text{:}\quad & -B_{t,T}(S_T - F_{t,T}) = -75 \times (0.80 - 0.96) &= +12.
\end{aligned}$$

From the table, we see that now not all uncertainty is gone: the deviations between cash flow and conditional expectations remain as large as before. That is because these deviations are the $\tilde{\epsilon}$'s, about which nothing can be done—at least not with currency forwards. But the conditional expected cash flows have been equalized, and as a result total risk is down. Again, this is the best reduction in the variance one can achieve, with these data.

General Minimum-Variance Hedging

When, realistically, the exchange rate can assume many more values than just two, it is generally the case that all conditional expected values no longer lie on a line. In fact, on the basis of our optimal-response argument we would expect cash flows to be convex in the exchange rate. Table 13.3 gives an example. It shows eleven possible exchange rates, the corresponding expected cashflows (in HC), and the probabilities of each. The slope of the regression is 87370, and the R^2 0.92. Figure 13.5 shows the original expectations for each exchange rate (the upward-sloping array of little squares); the regression line; and the hedged expected cash flows (the little triangles in a smile pattern).

Two remarks about these results, for the statistically initiated reader. First, note that since the data do not contain deviations from the conditional expectations, this is not the usual R^2: it tells you that the regression captures 92 percent of the variability of the conditional expected cash flows, not of the potential cash flows themselves. So this tells you that the non-linearity is not terrible, but you cannot conclude that hedging reduces risk by 92 percent since the residual risk is being ignored, here. Second, you may be wondering how the hedged-expectations series, which shows quite some curvature, still only contains just 8 percent of the variability of the original data. The answer is that the data are probability-weighted. The "distant" ends of the hedged series contain low-probability events that have only a minor impact on the variance. We are not used to this: our typical regression data in other applications are never weighted this way, or rather, we always let the sample frequencies proxy for the probabilties. Thus, our eye is trained to see each dot on the graph as equally probable, whereas here the central dots represent many observations. (In fact, the low-tech way to weigh the data is to repeat the observations such that their frequencies in the data matrix become proportional to the probabilities.) The weighting also explains why the regression line looks like mostly "below" the data. This is just because the regression line is heavily attracted by the central data, where most of the probability mass is.

Table 13.3: **Data for a non-linear exposure example**

S	0.80	0.82	0.84	0.86	0.88	0.90	0.92	0.94	0.96	0.98	1.00
V	42181	42821	43607	44572	45754	47203	48977	51148	53805	57054	61026
p	0.02	0.04	0.06	0.10	0.16	0.24	0.16	0.10	0.06	0.04	0.02

Figure 13.5: **Results for the non-linear exposure example**

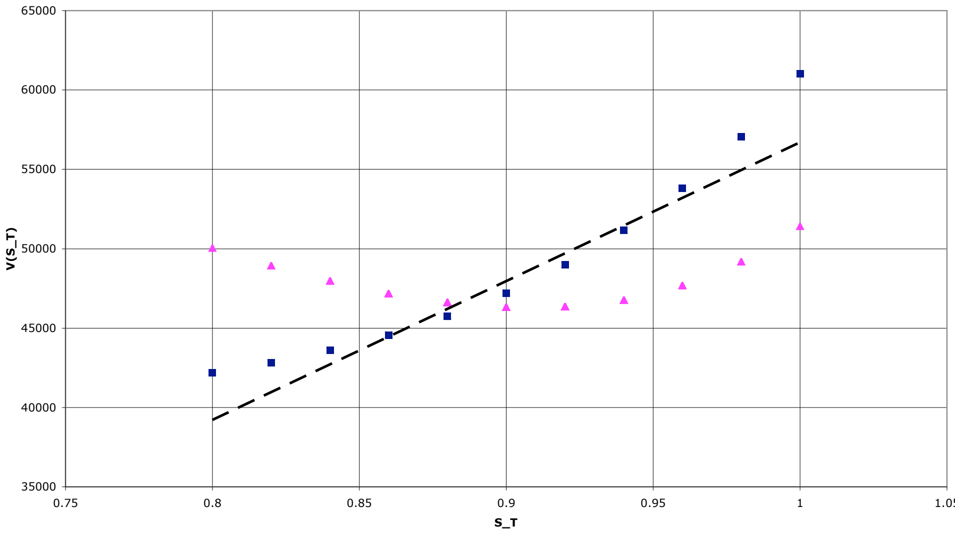

General Issues in Minimum-Variance Hedging

The above problems were kept simple, which is fine if the purpose is to explain the concept. Still, in fairness it must be added that the hedging of operations exposure is a bit of a minefield once you go to reality. Here is a list of the steps to be taken, and the issues to be solved:

- **Getting data** One can either go for data from the past, or numbers about possible future scenarios.
  - *Past data* One can proceed the way one estimates a market beta: collect past data on stock prices and exchange rates, and regress. We see the following problems. (*i*) This allows you, at best, to estimate the risk of the firm as a whole, not a new project or a separate business. (*ii*) The assumption is that the future is like the past, which is often not true: PPP deviations come in long swings, for instance, and exposure during a period of dollar overvaluation is a poor guide to exposure in the subsequent period of undervaluation. (*iii*) Even past exposure is estimated poorly because, for most firms, exchange risk is only a weak determinant of returns, which means that estimates are imprecise. (*iv*) If you nevertheless go for this data-mining approach, you should realise that, with time-series data, there is a problem of unit roots (ask your statistics teacher). This means that one has to use return data (percentage changes in values), not the value data themselves. The regression coefficient one gets from a returns regression is an elasticity (that is, $(\partial V/\partial S)(S/V)$) whereas the B we need is a partial derivative, $\partial V/\partial S$. So one would need an adjustment, multiplying the slope coefficient by V/S. Then a decision needs to

be made whether this correction will be based on the time series means, or the most recent values, or something else. We have no good guide to solve this issue.

– *Alternative scenarios for future cash-flows* The alternative to using time series of past data is to work with a cross section of alternative scenarios about the future. In principle this makes more sense. The only issue is the quality of these data in a real-world situation. The finance staff should know that the sales and cost data they get from Marketing and Operations, respectively, are crucial: if these are worthless, your hedge will be worthless too. Question them. Make sure the costs are not accounting COGS with mark-ups for overheads, for instance, but truly marginal cash outlays. Ask the marketing people how they would change their four or five P's under what scenario, thus forcing them to actually *think*.

- **Identifying the distribution of the future spot rate(s)** If you decide to work with scenarios, then you almost surely need to know how to weigh the possible pairs of possible future spot rates and associated expected cash flows. There are only three exceptions to this: weights are not needed if either it is reasonable to consider just two possible rates, like in the Android example, or if the expectations are linear in S, or if you go for a non-linear hedge (see below). But a two-point situation is exceptional; and if you get expected-cashflow data that are linear in S, that probably means the people who gave you the data were lazy: *a priori*, one expects convexity.

  Option traders typically start from a lognormal and then thicken the tails somewhat. You could get a standard deviation from them: ask for the "Implied Standard Deviation" or ISD. The mean, on the other hand, can be infered from the forward rate: we know there is, in principle, a risk correction that intervenes between CEQ()'s and E()'s, but it is small, both empirically and theoretically, and the choice of the mean has little impact on the regression anyway. Using forwards and ISDs, your forecasts for different horizons will be mutually compatible too. If you use a more wet-finger approach, compatibility over time is not guaranteed.

- **Linear or non-linear hedges?** When, realistically, the exchange rate can assume many more values than just two, it is generally the case that all conditional expected values no longer lie on a line. You then need to make up your mind as to whether you are happy with a static, linear hedge like we have discussed so far, or you prefer to go for a non-linear hedge. If, like in our example, the regression captures 92 percent of the expectations, you might be happy with the linear approximation and the associated hedge.

  The alternative is to go for a portfolio of options. In that case you construct a piecewise linear approximation to the data, using either your common sense (helped by pencil and ruler) or a regression with linear splines.[3] You start with a forward hedge whose size is, for instance, equal to the slope in the first linear section. With options you then let the exposure of your hedge portfolio change wherever you want, mirroring the changing exposures of your expectations. Alternatively, you can use dynamic replication of the options, but this introduces model risk: the dynamic replication will not do as well as the option itself, and how badly it deviates depends on the adequacy of the model chosen. Dynamic hedging is described in Chapter 9, on the Binomial Model.

The advantages of the non-linear hedge are twofold. First, you do not need the probability distribution of S: you leave this to the market, which then builds its perceptions about the density into the option prices. Second, there is a better fit with the data. The drawbacks include higher complexity, higher transaction costs, and perhaps over-reliance on expectations data that are more seat-of-the-pants than you may wish.

- **Hedging other risks?** If cash flows depend on other variables beside the exchange rate, and if for these other variables one also has forward or futures contracts, then you have the option to hedge the other exposures too. For instance, the oil price could be such a variable. We denote the additional variable by X, and there could of course be more than one extra X. The mean-variance hedge now requires that you run a multiple regression, $V = A + B \cdot S + C \cdot X$. For this, you'll need far more scenarios, and a joint probability distribution for X and S, which is not easy.[4]

13.3.3 Economic Exposure: CFO's Summary

Let us conclude this review of economic exposure by summarizing a few crucial results and integrating them with ideas mentioned in earlier chapters.

We can divide economic exposure into two categories—contractual exposure (aka

[3]First decide at what values of S you want a change of slope. These points are called knot points; for instance, in our example you may want a single change of slope, at $S = 0.90$ (right in the middle). Then make dummies $I_{k,j}$ indicating whether observation S_j is beyond the k-th knot point K_k; for instance, with one knot at $S = 0.90$, all observations with $S_j \geq 0.90$ get $I_{1,j} = 1$, and all lower observations get $I_{1,j} = 0$. Then regress $V_j = A + B_0 S_j + D_1[I_{1,j}(S_j - K_1)] + D_2[I_{2,j}(S_j - K_2)] \ldots$. The coefficient D_k tells you how much the slope changes in knot point K_k.

[4]Note that this makes sense only if you really want to hedge the additional risk with a linear hedge instrument, like oil futures or forwards. The econometrician's knee-jerk reaction is to add as many possible variables to a regression to improve the R^2 and isolate the contribution of S from that of other variables Z that are correlated with S. But if there is no hedge instrument for Z, sorting out the separate contributions of the two does not make sense. In fact, the difference between a multiple B and a simple-regression B is that the latter includes the effect of Z to the extent that Z resembles S. This is good, because then we at least do hedge the effect of Z (to the extent that Z resembles S).

transactions exposure) and operating exposure. Managers typically focus on contractual exposure, which arises from accounts receivable, accounts payable, long-term sales or purchase contracts, or financial positions expressed in foreign currency. This is because if one's source of information is accounting data, as it typically is, then transaction exposure is very visible and easy to measure. In contrast, operating exposure is much harder to quantify than contractual exposure; it requires a good understanding of competitive forces and of the macroeconomic environment in which the firm operates. For many firms, however, operating exposure is more important than contractual exposure, and it is critical that you make an attempt to identify and measure the exposure of operations to exchange rates.

Also, it is incorrect to assume that a firm with no foreign operations is not exposed to the exchange rate. For example, if a firm's competitors are located abroad, then changes in exchange rates will affect that firm's competitive position and its cash flows. Another common fallacy is the presumption that a policy of systematic hedging of all transaction exposure suffices to protect the firm against all exchange rate effects. As explained above, even if a firm perfectly hedges all contractual exposure, its operations are still exposed to the exchange rate.

Whether one considers transaction or operating exposure, one can use a forward contract (or the equivalent money-market hedge) to hedge the corresponding uncertainty in the firm's cash flow. Recall, however, that a forward or spot hedge is a double-edged sword. It is true that bad news about future operations is offset by gains on the forward hedge. However, you would likewise lose on the forward hedge if the exchange rate change improves the value of your operations. For example, in 1991, the Belgian group Acec Union Minière had hedged against a "further drop" of the USD. Instead, the USD rose, causing losses of no less than BEF 900m on the forward contracts. Four managers were fired. If you dislike this symmetry implicit in the payoff of a forward contract, you may consider hedging with options rather than forwards, to limit the downward risk without eliminating potential gains from exchange rate changes. As one banker once put it, "with a forward hedge you could end in the first row of the class or in the last; with an option, at worst you end somewhat below the middle."

A second potential problem that a treasurer needs to be aware of, when using short-term forward contracts to hedge long-term exposure, is the possibility of ruin risk, that is, liquidity problems that arise when there is a mismatch between the maturity of the underlying position and the hedging instrument. These liquidity concerns already came up in our discussion about hedging with futures contracts that are marked to market, but they arise any time the hedge triggers cash flows that come ahead of the exposed cash flow itself, for instance if a five-year exposure is covered by five consecutive one-year contracts.

Third, remember that, unlike many contractual exposures, operating exposure cannot be obtained from a balance sheet or a pro forma P&L statement. It has to be deduced from a cross-sectional analysis of possible future outcomes—cash flows, typically. The level of the true exposure can be totally out of the ballpark of the sizes of the exposed cashflows themselves, and can even have a different sign: remember the Android example.

Fourth, it is important to keep in mind that the estimate of exposure that one calculates changes over time, and may not be very precise at any given moment. However, this measure is useful—even if it gives us only an approximate indication of the sign and size of a firm's exposure—because it forces us to think about the way exchange rates affect the firm's operations.

Finally, hedging is like an aspirine—quite useful for short-term headaches but not a long-run remedy for most serious diseases. It does provide you with a financial gain that is intended to offset operating losses, but it does not reduce the operating losses themselves. One can live with operating losses as long as they are temporary; and the point of hedging, in such a case, is that it does provide the cash that tides you over a bad patch. But if the problem is likely to be more than just temporary, you need strategic changes in operations—for example, revising the marketing mix, reallocating production, choosing new sourcing policies to reduce exposure, and so on. Again, financial hedging just provides cash that eases the pain and helps financing the adjustments; it does not solve the underlying problem.

In this respect, when making scenario projections about the possible future exchange rates, we should also make contingency plans for various possible future exchange rates, including less likely ones. One can win crucial time if the response has been talked through before; otherwise one wastes too much time deciding what exchange-rate changes are "big" and "structural" or not, what the available options are, not to mention who should be on the "task force" that ruminates on all this, and so on and so forth.

This finishes our discussion of economic exposures. We now turn to translation exposure.

13.4 Accounting Exposure

The $\tilde{V}_T$ entry in the exposure definition has been interpreted, thus far, as the portfolio of contractual FC-denominated undertakings inherited from the past, or a portfolio of activities that need continuous decisions influenced by, amongst other things, exchange rates. The third definition we discuss is the firm's accounting value. This accounting value may be affected by exchange rates in two ways. First, the firm may have contractual exposures which the firm is also *marking to market*, thus adjusting their book values to the rates that prevail on the valuation date. Second, the firm may have foreign subsidiaries, and the HC value of their *net worth*, in the accounting sense of the word, probably depends to some extent on the exchange rate that prevails on the consolidation date.

13.4.1 Accounting Exposure of Contractual Forex Positions

The issue of how to book contractual exposures has been brought up already in Chapter 5, where we argued that translation at the then prevailing forward rates makes more sense. Still, many firms use the spot rate. The issue here is different, though. Notably, if the firm has booked a contractual position in the past, should it adjust the book value on the reporting date, and, if so, how?

A/P, A/R, deposits, loans For these items, both US GAAP and the IFRS rules would agree, sensibly, that marking to market is recommendable; in the case of IFRS, that even is the general rule. Our earlier logic would then imply that the current forward rate be used to translate the values of A/R, A/P, deposits, loans, etc. into HC. (Ideally, one would also correct for time value and changes therein by PV'ing all numbers, but this is still too rarely done even though IFRS supports this). Any increase of the value of an asset would be balanced by an increased "liability", the unrealized capital gain that adds to the shareholders' net worth. (Similar statements can be made for losses, and for short positions, of course.) Being unrealized, many managers would prefer that the gain would not pass through Profit&Loss first, but IFRS begs to differ.

Futures For futures hedges and the like, the same logic holds. Instead of mentioning a zero value off balance sheet for a futures contract, one can add a capital gain or loss $f_{t,T} - f_{t_0,T}$. This entry is the counterpart, on the liability side, of all net marking-to-market cash flows that have been received from the Clearing Corporation since the initial value date t_0 (or the beginning of the accounting period if there has been at least one earlier financial report, which presumably contains the gains/losses prior to that date) and, therefore, have already shown up in "bank account", on the assets side of the balance sheet. If the marking-to-market cash flows have the character of a final payment rather than adjustments to security posted—the tell-tale symptom would be that there is no interest earned on outward payments, nor due on inward payments—then one could argue that the gain or loss is realised and, therefore, should be shown as part of P&L rather than just as an unrealized item among the shareholders' funds. This is the FASB position. The IAS position, as reflected in IFRS, is that all gains or losses have to be shown, whether realized or not.

Note that, if the firm has taken out a futures contract to hedge another position, and if that other position is not being marked to market and if the firm has to book its marking-to-market cash flows on the futures as a profit or loss, then realized profits become more volatile even though the hedging aims to reduce variability. In that case, the FASB would waive the requirement to book the gains and losses via P&L, provided the futures position was immediately designated as a hedge of a well-identified balance-sheet item. There is no such rule for forwards (where, by FASB rules, marking-to-market does not have to go through P&L) or cash hedges (where, presumably, the firm's marking-to-market rules for hedge and hedgee are always in agreement). But there is no similar rule either for exposures that are not

Table 13.4: Valuation using IFRS, pure spot, and pure forward rules

		balance-sheet items						gains (+) or losses (-) on marking to market of ...		
		Cash & bank	Invntry	A/P	*budgeted asset*	*budgeted liability*	forward	cmmtmnt	A/P	forward
CURRENT IFRS										
15-Oct S=1.000; F=1.018	order	n.a.	n.a.	n.a.	*1,018*	*1,018*	0		n.a.	
15-Nov S=1.020; F=1.035	book invoice a S			1,020		*1,018*		-2		
	book goods at S-dF		1,003		*-1,018*			-15		
	marking to market of hedge						17			17
	subtotals		1,003	1,020	*0*	*0*	17	-17		17
31-Dec S=1.040; F=1.045	marking to market of A/P			20					-20	
	marking to market of hedge						10			10
	subtotals		1,003	1,040	*0*	*0*	27	-17	-20	27
15-Jan S=1.015=F	marking to market of A/P			-25					25	
	marking to market of hedge						-30			-30
	pay bill, close A/P \& forw	-1,018		-1,015			3			
	subtotals	-1,018	1,003	0	*0*	*0*	0	-17	5	-3

Key The entries shown in italics are not actually used in IFRS but help explain what is done. A good is ordered at t_o for a foreign-currency price 1,000. At the time of ordering, the firm is assumed to record the future goods and the future A/P into its budget at the initial forward rate, 1.018. It closes these budget accounts and makes genuine accounting entries when the goods are actually delivered. At t_b, the invoice is entered at spot (1.020), leaving a loss of 2 relative to the initial valuation of 1.018. At t_b, the interim gain on the forward purchase is recognized (+17 pips). The cost of of the good is recognized to be only 1003, namely, the A/P value 1.020 minus the gain on the hedge, 17. The total result on the commitment consists of –2 (when the liability is entered as 1.020 instead of the initial entry, 1.018) and –15 when the asset is booked at 1.003 instead of its initial valuation, 1.018.

At year end, all positions are marked to market. The bookings for the date the invoice is paid start by marking everything to market again and then realizing the total loss on the forward contract (3, when currency worth 1.015 is bought at the original forward price of 1.018. The three lines can. of course, be merged into two or even one line.

yet in the balance sheet, which is anomalous, economically. Also, while the rule says that "speculative" futures positions should be fully marked to market, there is no such requirement for speculative positions in forward or spot markets.

Forwards For forwards there is no cash movement prior to expiry, so the accounting entries in case of a gain on a long position would be (i) a revaluation of an asset with original book value zero, and (ii) an upward adjustment in shareholders' funds, possibly as an unrealized and undistributable item. Again, almost surely the time value part that we showed in Chapter 4 would be missing: only the un-PV'ed part $F_{t,T} - F_{T_0,T}$ would be reported.

IFRS prescribes that all forward positions be shown—initially at zero value, and later marked to market using the change in the forward rate (undiscounted). The A/R or A/P position is to be booked at the spot rate, and marked to market at the spot rate. So the marking-to-market (M2M) of hedge and hedgee will roughly match but the difference between initial spot and forward is treated as a capital gain or loss—a bad idea, I argued in Chapter 5, because laypersons will think it actually *means* something.

The rule that forwards need to be M2M-ed creates a problem if the hedge is undertaken before the invoice is written or received: then, pending the invoice, the forward contract would already trigger M2M cashflows while there are no roughly-matching M2Ms for a hedgee yet. Thus, again, the hedge would add uncertainty to the reported results even though it actually stabilizes cash flows. To solve this, IFRS could have decided to start booking the hedgee transaction at the date of *firm commitment*—presumably also the date the hedge is undertaken—rather than the date of the invoice or transfer of ownership, but that would have been a major change in accounting practices. For this reason, IFRS concocts an account 'firm commitments', which just absorbs any gains and losses in the forward contract during this initial period. For this to be possible, the hedge must be immediately designed as such, and linked to the specific transaction.

Example 13.8

Kayblan Whyer US orders glass fiber cables worth EUR 1m on October 15. The wares are shipped and invoiced on November 15. On December 31 the A/P is marked to market, and it is paid on January 15. In Table 13.4 I show the entries—as numbers in columns that stand for accounts, not as debits and credits, because columns are easier to follow. The columns on forwards and firm commitments are relevant only if Kayblan does hedge.

* * *

This finishes our brief discussion on updating the book values of contractual exposures. But these are only part of the balance-sheet items that might be affected. As mentioned already, also a subsidiary's P&L and A&L statements may have to be translated, for consolidation of the accounts, for instance. We first list some of the reasons why the financial statements of a subsidiary need to be translated into the currency of the parent firm. Next we describe the four main translation methods. We conclude with a discussion of the relevance of translation exposure.

13.4.2 Why Firms Need to Translate Financial Statements

If some of the subsidiaries of a firm are located abroad, their financial statements are typically maintained in terms of the local currency, which is foreign to the parent. There are a number of reasons why the financial statements of the subsidiary may need to be translated into other currencies—most often, the parent company's home currency:

- **Taxes** Translation is often necessary for tax purposes, notably if the tax authorities of the parent's home country have to review the subsidiaries' financial statements to establish the tax basis (as explained in Chapter 20). Taxes in the parent's home country, on income earned by the foreign subsidiary are, of course, payable in home currency. This means that the foreign income has to be translated into the home currency. Also, capital gains arising from exchange rate changes may be taxable; if so, to compute the capital gain, one needs to translate the value of the foreign subsidiary into home currency terms. Thus, translation exposure, even though it deals with accounting data, can have an impact on cash flows through its effect on the tax basis.

- **Consolidated Financial Statements** Most countries require consolidation of the parent's and subsidiaries' financial statements for reporting purposes. Consolidation here refers to the integration of the financial statements of the firm's subsidiaries into the parent's asset and liabilities (A&L) and profit and loss (P&L) statements. Of course, one needs to first translate the financial statements of the subsidiary before they can be consolidated with those of the parent.

- **Performance evaluation and budget allocation across subsidiaries** The parent firm itself may feel the need to translate the financial statements of foreign subsidiaries. This is because one needs to compare data in order to allocate investment budgets or to evaluate the performance of the subsidiary. For example, even to get some idea about the importance of a foreign unit, one needs to determine its value in terms of a common currency. Of course, importance cannot be determined on the basis of a single figure and surely not on the basis of just backward-looking accounting data. Still, translated accounting data give a first impression of the relative importance of the foreign activities.

- **Bonuses** In order to make performance measures comparable, foreign data need to be translated into a common currency. For example, many firms have bonus plans that link their managers' compensation to their performance. Decisions to promote or fire managers are also based on performance. To make such decisions, one needs to translate the financial statements of the foreign subsidiaries into the currency of the parent.

- **Valuation** To value the entire firm (as an outside investor or financial analyst), one needs far more than just accounting data. Still, valuation is often partially based on accounting values; or, at the very least, the accounting value serves as a benchmark. For instance, if the discounted cash flow value of the entire firm turns out to be four times its book value, one would surely take a closer look at both types of information. Again, the book value of the firm as a whole cannot be computed unless assets and liabilities of foreign subsidiaries are first translated into a common currency.

In the next section we first discuss the general objectives that any method used to translate the accounts of the subsidiary into the currency of the parent firm tries to accomplish, and then the details of the various methods that are used for translation.

13.4.3 The Choice of Different Translation Methods

Accounting exposure arises because the outcome of translating a subsidiary's balance sheet from foreign currency to home currency depends on the exchange rate at the date of consolidation, an exchange rate that is uncertain. Firms may like to hedge this exposure to reduce or eliminate the swings in reported profits that arise simply due to these translation effects. This exposure, of course, depends on the rules used to translate the accounts of the subsidiary into the currency of the parent firm. There are a variety of approaches that one can adopt to translate the income statement and balance sheet items of the subsidiary into the currency of the parent firm.

Example 13.9

Suppose a Canadian firm buys a competitor in England for GBP 1m, when the exchange rate is CAD/GBP 2.0. A year later, the exchange rate is CAD/GBP 2.1. Thus, assuming that the subsidiary is still worth GBP 1m and translation is done at CAD/GBP 2.1, its translated value in terms of the currency of the parent is CAD 2.1m. One question is whether one *should* translate the GBP value at the new rate at all; and, if the answer is positive, the next question is how to report this increase in the value of the British subsidiary in the accounts of the parent firm. For example, should the exchange rate effect be shown as part of the reporting period's income, or should it just be mentioned on the balance sheet, as an unrealized gain?

If the decision is to translate at the historical exchange rate—the one prevailing when the asset was purchased—then there is no translation exposure. Otherwise there is, but its size depends on how one translates; for example, one could opine that real assets do not really lose value following a devaluation, etc.

The above example illustrates what the controversy between accountants is all about. Accountants do not agree which assets and liabilities should be translated at the historical exchange rate and which at the "current" or "closing" exchange rate, that is, at the rate prevailing at the date of consolidation. There is also some disagreement about whether and when exchange rate gains or losses should be recognized in income. A major criterion of accountants in devising the translation rules is whether these rules are consistent with the rules for domestic accounting. However, from a firm's point of view, the principal requirement is that the rules be such that they provide accurate information about the performance of the subsidiary. Lastly, firms also wish that the rules be such that they do not lead to wide swings in the figures reported in the financial statements.

In the rest of this section, we describe four different translation methods and the philosophy underlying each method. Each method has a set of rules for translating items in the balance sheet and the income statement. The rules for translating items in the income statement are quite similar across the different methods; hence, we will focus on the rules for items reported in the balance sheet. To illustrate the differences between these methods, we shall consider the example of an Australian

Table 13.5: **Translating the Australian Balance Sheet into MTL**

	Value in AUD	values in MTL after translation at 0.333 or 0.300: curr/noncurr (at .333)	curr/noncurr (at .3)	mon/nonmon (at .333)	mon/nonmon (at .3)	closing rate (at .333)	closing rate (at .3)
Assets							
cash, securities	1,000	333	300	333	300	333	300
A/R	1,000	333	300	333	300	333	300
inventory	1,000	333	300	*325*	*325*	333	300
plant, equipment	5,000	*1,625*	*1,625*	*1,625*	*1,625*	1,665	1,500
Total assets (a)	8,000	2,624	2,525	2,616	2,550	2,664	2,400
Liabilities							
A/P	500	166.5	150	166.5	150	166.5	150
Short-term debt	2,000	666.0	600	666.0	600	666.0	600
Long-term debt	2,400	*780.0*	*780.0*	799.2	720.0	799.2	720.0
Total Debt (b)	4,900	1,612.5	1,530	1,631.7	1,470	1,631.7	1,470
Net worth (a)–(b)	3,100	1,011.5	995.0	984.3	1,080.0	1032.3	930.0
of which:							
Retained	0	*0*	*0*	*0*	*0*	*0*	*0*
Equity	3,100	*1002.0*	*1002.0*	*1002.0*	*1002.0*	*1002.0*	*1002.0*
Eq Adjustment	—	9.5	7.0	17.7	78.0	30.3	72.0
exposure of net worth		$\frac{1011.5-955.0}{0.333-0.300}$ = AUD 500 = nt wrkng cptl		$\frac{984.3-1080.0}{0.333-0.300}$ = AUD–2,900 = nt mon. assets		$\frac{1032.3-930.0}{0.333-0.300}$ = AUD 3.100 = net worth	

Key Items deemed to be non-exposed are printed in italics.

subsidiary of a Maltese firm. A simplified balance sheet of the subsidiary is shown in the second column (value in AUD) of Table 13.5. We shall explain the notion of accounting exposure by considering translation on December 31, 2007, at two different exchange rates, MTL/AUD 0.333 and MTL/AUD 0.300, and by seeing how the value of the subsidiary changes depending on the accounting method being used. Throughout this discussion, our focus will be to study what the different translation methods imply for the firm's accounting exposure.

The four methods all share the following steps: (*i*) translate assets and debt, using the method's rules as to what items are exposed or not; (*ii*) compute net worth (assets minus debts, in HC); (*iii*) subtract equity at historic valuation (including past retained earnings, each at its own historic valuation) to identify the balancing item, Equity Adjustments.

The Current/Non-Current Method

The Current/Non-Current Method for translating the financial statements of foreign subsidiaries is one that was commonly used in the US until the mid-1970s. As its name suggests, whether an item is translated at the closing exchange rate or the historical rate depends on its time to maturity. Thus, according to this method, current (*i.e.* short-term) assets and liabilities in the balance sheet are translated at the closing exchange rate, while non-current items, such as long-term debt, are translated at the historical rate. The logic underlying this is that the value of short-term assets and liabilities is fixed, or at least quite sticky, in AUD terms, so that its HC value changes proportionally with the exchange rate. For example, the future value of a AUD T-bill is fixed in AUD nominal terms; and, in the short-term, goods prices are sticky and therefore quasi-fixed in AUD terms, too. Long-term assets and liabilities, in contrast, will not be realized in the short run—and by the time they are realized, the closing exchange rate change may very well turn out to have been undone by later, opposite changes in the spot rate. That is, the effect of a closing exchange rate change on the realization value of long-term assets and liabilities is very uncertain. As accountants hesitate to recognize gains or losses that are very uncertain, the Current/Non-Current Method simply prefers to classify the long-term assets and liabilities as unexposed.

Thus, under the Current/Non-Current Method, translation at the closing rate is restricted to only the short-term assets and liabilities. Thus, exposure is given by the difference in short-term assets and liabilities, that is, Net Working Capital.

Example 13.10

In Table 13.5, we assume that long-term debt was issued and long-term assets (plant and equipment) were bought in early 2007, at which time the exchange rate was MTL/AUD 0.325. Thus, these items are recorded at their historical values (indicated as italicized text) and are not affected by the exchange rate. It follows that net exposure equals short-term assets minus short-term liabilities, or net working capital—AUD 500. The effect of the exchange-rate change from 0333 to 0.300 is a drop in net worth of AUD $500 \times (-0.033) = -$MTL 16.5.

Evaluation

- The assumption underlying this method seems to be that there is mean-reversion in exchange rates; that is, exchange rate fluctuations tend to be undone in the medium run, which (if true) means that they affect short-term assets only. However, as discussed in Chapter 11, there is little empirical support for this view (except for the small movements of exchange rates around a central parity): typically, changes in exchange rates are not undone in the medium run, and floating exchange rates behave like random walks.

- Most firms have positive net working capital and would therefore be deemed to be positively exposed (losing value, that is, following a devaluation of the host currency). Yet economic logic says that the true effect on economic value should be hard to predict in general, depending, for instance, on whether the firm is an exporter or an importer, a price taker or a price leader, in a small open economy or in a large, closed one, or competing against locals or versus foreign companies, etc. Thus, there is little hope that this method will capture the true value effect except by pure serendipity.

- The consolidated accounts are not compatible with the subsidiary's original accounts. The relative values of items differ according to whether one uses HC or FC numbers, and many of the standard ratios will be affected. This is not good news if e.g. performance analysis is based on ratios.
- The resulting translated balance sheet is a mixture of actual and historic rates and, therefore, hard to interpret.

To translate the subsidiary's income statement, the Current/Non-Current Rate Method uses an average exchange rate for the period, assuming that cash flows come evenly over the period—except for incomes or costs corresponding to non-recurrent items (like depreciation of assets): these are translated at the same rate as the corresponding balance sheet item. This creates another inconsistency between the AUD and MTL P&L figures, and between the translated P&L and A&L figures.

The Monetary/Non-monetary Methods

The Monetary/Non-Monetary Method and its close kin, the Temporal Method are said to be ideally suited if the foreign operation forms an integral part of the parent. The idea is that, accordingly, the translation should stay as close as possible to what would have happened if the operation had been run as a branch, that is, just a part of the main company that happens to be active abroad and has assets abroad but does not have a separate legal personality.

If the foreign business had been a branch indeed, without any separate accounting system, the translation issue would not have arisen: everything would have been in the parents' books already, in HC, except for monetary assets whose value by definition is fixed in FC terms and needs to be translated. For instance, if the parent firm held forex cash or other monetary assets expressed in forex, any value change would have been recorded and probably included into the parent's P&L; but its machines and buildings would have been unaffected, in terms of book value, by exchange-rate changes. Since by assumption the subsidiary is really a part of the parent, the subsidiary's monetary A&L are translated at the closing rate, and the non-monetary items at the historic rate. Any resulting gains or losses are mentioned among the reserves, as unrealized gains or losses.

The above argument assumes that domestic assets are valued at historic cost, which principle is becoming less and less popular. But there exist another angle to justify the rule. It is sometimes argued that, in the long run, inflation differentials should undo exchange rate changes (PPP). So in the long run the real value of real assets will not be affected. Thus, according to this method, we should adjust only the monetary (not the real) assets and liabilities for changes in the exchange rate. It follows that only the net foreign-currency monetary position, financial assets minus debt, is exposed.

Example 13.11

In Table 13.5, the Net Worth figures under each of the two possible year-end exchange rates differ by MTL +95.7. The exposure was $2000 - 4900 = -2900$ (AUD), under this method.

Evaluation

- The Purchasing Power Parity view of the world has received little empirical support, except vaguely in the never-arriving long run.[5] Accountants usually do not rely on highly uncertain prospects.
- In addition, PPP just says that translated values of assets abroad tend to be equal to values at home. If true, this would mean that changes in values of foreign and domestic assets are equal to each other; but there is no claim that the value change at home is zero.
- Likewise, in the "closely related operations" versions of the story the non-monetary assets are treated like they would have been treated if they were at home and, therefore, left unchanged. But that's historic costing. In many cases, under replacement value or fair value the value of the foreign-based asset would differ from one at home, and the argument would break down.
- This measure of exposure, financial assets minus debts, is likely to be negative for most firms. Thus, under the Monetary/Non-Monetary Method, a devaluation will typically lead to an accounting gain rather than to a loss. But, from our earlier discussion on a related point in the Current/Noncurrent method, economic reality should be very different for different firms, so the hope that this method produces the true number is, again, slim.
- Also here, HC relative values differ from FC ones, affecting ratios; there is no consistency.
- Finally, the resulting mixture of actual and historic translations is again hard to interpret.

To translate the subsidiary's income statement, the Monetary/Non-Monetary Method uses an average exchange rate for the period, except for incomes or costs

[5] Recall that all we know is that uncertainty about future real exchange rates does not grow proportionally with the length of the horizon, which is a far cry from uncertainty somehow disappearing entirely the longer one waits.

corresponding from non-monetary sources (like depreciation of assets). These are translated at the same rate as the corresponding balance sheet item. This again creates an inconsistency between the AUD and MTL P&L figures, and between the translated P&L and A&L figures.

The Temporal Method

The Temporal Method of translating the financial statements of a foreign subsidiary is similar to the Monetary/Non-Monetary Method. One difference between the two methods arises if "real" items have been marked to market in HC. As we saw, under the monetary system, inventory is always translated at the historical exchange rate, since it is a non-monetary asset. Under the Temporal Method, in contrast, inventory may be translated at the current (*i.e.* closing) exchange rate if it is recorded in the balance sheet at current market prices. The advantage of this approach is that it is less inconsistent with the accounting rules used for the parent firm if the parent marks to market its domestic inventory too. Another aspect of the temporal method is that it makes translation effects part of reported income, which can lead to large swings in reported earnings. Thus, under this method CFOs tried to hedge exposures that were just arbitrary paper results, not real cash flows.

The US accounting directive that was used from 1976 to 1981, FASB 8, was based on the Temporal Method. (Prior to that, the US imposed the Current-Noncurrent method.) The Closing Rate Method, introduced by FASB 52 in the US, is designed to overcome some of these problems.

The Current- or Closing-Rate Method

This is the simplest approach for translating financial statements. According to the Current Rate Method or Closing Rate Method, all balance sheet items are translated at the closing exchange rate. Typically, exchange gains are reported separately in a special equity account on the parent's balance sheet, thus avoiding large variations in reported earnings, and these unrealized exchange gains are not taxed.

Example 13.12

For the Australian subsidiary's simplified balance sheet, the exposed amount is net worth, AUD 3,100.

Evaluation

- The main advantage is consistency between the parent's and subsidiary's relative numbers. Likewise, using one rate produces a number that is easier to interpret than one resulting from mixtures of closing and historic translations.[6]

[6] If the subsidiary's accounts themselves mix historic costs—some of them possibly very dated—

- Under this method, a 15% devaluation means a 15% decrease in the net value of the investment. Economically, one expects that a devaluation of, say, 15 percent leads to a value loss of 15 percent only if all subsequent cash flows are unaffected (in HC terms), which assumes a very closed economy. So, again, this method is unlikely to capture the true economic effect.

To translate the income statement, one translates all items at either the closing exchange rate or the average exchange rate of the reporting period. The first method is chosen for consistency with the balance sheet translation. The second method is based on the argument that expenses that have been made gradually over the year should be translated at the average exchange rate. (Curiously, this argument does not seem to apply to expenses that end up capitalized into balance-sheet items.) Profits, the argument goes, are realized gradually over the year, and should be translated at an average rate. This, of course, contradicts the translation of the balance sheet at a single exchange rate.

13.4.4 Accounting Exposure: CFO's Summary

As we have seen, there are various methods to translate a subsidiary's balance sheet into the parent's currency. Many regulating bodies favor the Closing Rate Method. For example, the US Financial Accounting Standards Board has essentially imposed this method (FASB #52, 1982) for most operations, and allows the old Temporal Method only for foreign operations closely integrated with the domestic headquarters. Similar rules were issued soon thereafter in the UK and Canada. The International Accounting Standards Committee has likewise come out in favor of the Closing Rate Method (IASC #21, 1983—a text that, unlike FASB #52, is well-written, lucid, and short), again except for closely related operations, where the Temporal Method is imposed.

However, the IASC can provide recommendations only; it has no statutory power to impose accounting rules anywhere. In continental Europe there is no consensus as to what method is to be followed. For example, in many countries (including, until the early nineties, Italy and Belgium), consolidation was not mandatory and, therefore, not regulated, while in other countries (including Germany), the obligation to consolidate was not extended to foreign subsidiaries. The EC 7th Directive, passed in 1983 and implemented in most member states by the early nineties, imposes consolidation but does not prescribe any particular translation method. The only requirement is that the notes to the accounts should disclose the method that was used. Only under IFRS, the IAS rules do apply; but in the EU IFRS is mandatory only for listed companies and financials. Other companies can use traditional local GAAP, which typically leaves considerable discretion.

---

with true(er) recent valuations, the result remains hard to interpret. But at least the translation procedure no longer adds to that problem.

Given the wide choice that is offered in many cases, one could wonder which method is best. And even where one particular method is imposed, one could consider whether it is useful to adopt a different method for internal reporting purposes. Even more fundamentally, one could ask whether accounting exposure matters at all. Let us briefly dwell on this before we close this chapter.

From the discussion of the various translation methods, we see that the question of which translation method to choose is similar to the issue of whether the firm should use the method of last-in/first-out (LIFO), or first-in/first-out (FIFO), or some average cost, for the purpose of valuing its inventory. One could argue that the accounting method for inventory valuation does not matter, since a shift from, say, LIFO to FIFO will change neither the firm's physical inventory nor its cash flows (except possibly through an effect on taxes). Moreover, one could argue that neither LIFO nor FIFO nor average cost is correct; only replacement value is theoretically sound. In the same vein, one could argue that the choice of the translation method does not affect reality—except possibly through its effect on taxable profit—so that the whole issue is, basically, a non-issue. Furthermore, while in the case of inventory valuation, one could argue that LIFO, being generally closer to replacement value, is the least of all evils, it is not obvious which of the translation methods generally corresponds best to economic value. The whole issue is, perhaps, best settled on the basis of practical arguments. Accounting data are already complicated enough, so that the Current Rate Method is probably a good choice, given its simplicity and internal consistency.

In Table 13.6 we compare economic and accounting exposures. Perusal of the list will reveal that economic exposure is the one to watch, not accounting exposure. But although accounting exposure suffers from the limitations described above, often accounting data are the only data that are readily available to a firm. Thus, it is important that treasurers and CFO's, when using these data to make hedging decisions, be aware of these limitations when using accounting data.

* * *

Let is recapitulate the results obtained thus far in the current part. We first argued that hedging adds value at least for some firms some of the time. We then discussed exposure, that is, the size of the forward hedge that should be added to minimize uncertainty. But the decision whether or not to hedge was hitherto discussed in isolation from other risks the firm incurs, many of which are not hedgeable at all. So perhaps the question should be what the total risk of the company is, and by how much this total risk goes down if exchange risk is being hedged. Such a holistic, portfolio view is taken by Value at Risk (VaR), the issue of the next chapter.

Table 13.6: **Economic v Translation Exposure: summary**

1. Economic exposure relates to changes in genuine cash flows and their PV.	Accounting exposure focuses on book values with, usually, no cashflow repercussions. (One possible exception is through taxes, if translation gains are taxed.)
2. Economic exposure is forward looking: it relates to future cash flows.	Accounting exposure is backward looking: it relates to past decisions on assets and liabilities as recorded on the balance sheet.
3. Economic exposure covers all cash flows, whether or not they can be found in the current financial statements.	Accounting exposure is confined to A&L and P&L items.
4. Economic exposure exists for virtually all firms.	Accounting exposure only exists when there are FC-denominated A&L items or subsidiaries whose accounts need to be consolidated.
5. Economic exposure depends on economic facts, like the contracts the firm signed or the economic environment it operates in.	Accounting exposure depends on the translation method chosen or prescribed, without reference to the economic framework.

13.5 Test Your Understanding: contractual exposure

13.5.1 Quiz Questions

True-False Questions

1. Exchange risk describes how volatile a firm's cash flows are with respect to a particular exchange rate.
2. Exchange exposure is a measure of the sensitivity of a firm's cash flows to a change in the spot exchange rate.
3. Hedging exposure means eliminating all risk from a net position in a foreign currency.
4. If you need to hedge a series of exposures with different maturities and you use duration hedging, it is best to hedge the negative exposures separately from the positive exposures.
5. Contractual exposure is the absolute change in the firm's cash flows for a unit change in the spot exchange rate.

6. Operating exposure is the exposure that results when the forward rate is at a discount with respect to the spot rate at the moment you sign a sales or purchase contract.

7. Contractual exposure is additive for one maturity and one currency.

8. Options are undoubtedly the best choice for hedging foreign currency exposure because the possibility of profiting from a favorable change in the exchange rate remains open without the losses from an unfavorable change in the exchange rate.

9. Reverse exchange risk is the risk that arises when you receive a foreign currency A/R that you left unhedged, and the exchange rate at the time of receipt is unexpectedly low.

10. When interest rates are zero, we can aggregate exposures of a given currency across time.

11. If interest rates are positive but certain, and exchange rates are uncertain, we can aggregate the exposure of one currency across time once we take time value into account.

12. By pooling the aggregate exposure of one currency across time, we can ignore time value, because we have arbitraged away interest rate risk. The only risk that remains is exchange rate risk.

Matching Questions

Suppose that you are a manager at a British firm, and you are responsible for managing exchange rate exposure. Determine whether the following statements are related to accounting exposure, operating exposure, or contractual exposure.

1. Your German subsidiary has recently made new investments.

2. You bought a call option on EUR to hedge an EUR accounts payable.

3. You have just sold goods to an American customer. The customer has ninety days to pay in USD.

4. You have just developed an exciting new product. The success of this product depends on how it is priced in the local currencies of your export markets.

5. You have made a bid to deliver your exciting new product to schools in France during the next academic year. You will learn whether or not the bid has been accepted in three months.

6. You sell wool but face potential competition from Australia. If there are no imports, the price of your wool will be GBP 1. However, Australians enter your market once the exchange rate falls below GBP/AUD 2.

13.5.2 Applications

1. The American firm, American African Concepts, has a one-year EUR A/P totaling EUR 100,000 and a one-year Senegalese A/R totaling CFA 120,000,000. The CFA/EUR exchange rate is fixed at 655.957.

    (a) Can AAC offset its EUR A/P with its CFA A/R?

    (b) If so, how much exposure remains?

2. The Dutch manufacturer Cloghopper has the following JPY commitments:

    - A/R of JPY 1,000,000 for thirty days.
    - A/R of JPY 500,000 for ninety days.
    - Sales contract (twelve months) of JPY 30,000,000.
    - A forward sales contract of JPY 500,000 for ninety days.
    - A deposit that at maturity, in three months, pays JPY 500,000.
    - A loan for which Cloghopper will owe JPY 8,000,000 in six months.
    - A/P of JPY 1,000,000 for thirty days.
    - A forward sales contract for JPY 10,000,000 for twelve months.
    - A/P of JPY 3,000,000 for six months.

    (a) What is Cloghopper's net exposure for each maturity?

    (b) How would Cloghopper hedge the exposure for each maturity on the forward market?

    (c) Assume that the interest rate is 5 percent (compound, per annum) for all maturities and that this rate will remain 5 percent with certainty for the next twelve months. Also, ignore bid-ask spreads in the money market. How would the company hedge its exposure on the spot market and the JPY money market? Describe all money-market transactions in detail.

    (d) If the interest rate is 5 percent (compound, per annum) for all maturities and will remain 5 percent with certainty for the next twelve months, how would the company hedge its exposure on the forward market if only one forward contract is used?

    (e) Assume that Cloghopper prefers to use traded options rather than forward contracts. The option contracts are not divisible, have a life of either 90, 180, 270, or 360 days, and for each maturity the face value of a contract is JPY 1,000,000. How could Cloghopper hedge its exposure? Do the options offer a perfect hedge for each maturity?

    (f) Drop the assumption of a flat and constant term structure. If Cloghopper wants to hedge its exchange rate exposure using one forward contract and its interest rate exposure using FRA contracts, how would the analysis of parts (c) and (d) be affected? A verbal discussion suffices.

(g) The term structure is flat right now (at 5 percent p.a., compound), but is uncertain in the future. Consider the spot hedge of part (c). If, instead of FRAs, duration is used to eliminate the interest risk, how should Cloghopper proceed?

13.6 Test Your Understanding: Operating exposure

13.6.1 Quiz Questions

True-False Questions

1. A firm that has no operations abroad does not face any operating exposure.
2. Only firms with exports, or firms that compete against foreign exporters, face operating exposure.
3. A firm that denominates all of its contracts in home currency, or hedges all of its foreign currency contracts, faces no operating exposure.
4. Almost every firm faces some operating exposure, although some firms are only exposed indirectly (through the country's general economic activity).
5. As large economies have a big impact on world economic activity, companies in such countries tend to be very exposed to exchange rates.
6. Small economies tend to fix their exchange rate relative to the currency of larger economies, or tend to create currency zones (like the EMS). Therefore, companies in small economies tend to be less exposed to exchange rates.
7. The smaller a country, the more open the economy. Therefore, exposure is relevant for most of the country's firms.
8. Everything else being the same, the larger the monopolistic power of a firm, the smaller its exposure because such a firm has more degrees of freedom in adjusting its marketing policy.
9. Consider an exporting firm that has substantial monopolistic power in its product market. Everything else being the same, the more elastic foreign demand is, the more an exporting firm will profit from a devaluation of its own currency. Similarly, the less elastic foreign demand is, the less an exporting firm will be hurt by an appreciation of its own currency.
10. Most information needed to measure operating exposure can be inferred from the firm's past export and import contracts.

Multiple-Choice Questions

Choose the correct answer(s).

1. In a small, completely open economy,

   (a) PPP holds relative to the surrounding countries.

(b) A 10 percent devaluation of the host currency will be offset by a 10 percent rise in the host country prices.

(c) The value of a foreign subsidiary, in units of the foreign parent's home currency, is unaffected by exchange rate changes.

(d) The real value of a foreign subsidiary to an investor from the host country is unaffected by exchange rate changes.

(e) In the absence of contracts with a value fixed in the host currency, the real value of a foreign subsidiary to an investor from the parent's home country is unaffected by exchange rate changes.

(f) In the absence of contracts with a value that is fixed in foreign currency, the real value of a foreign subsidiary to an investor from the host country is unaffected by exchange rate changes.

(g) There is little or no advantage to having one's own currency: exchange rate policy has virtually no effects.

2. In a completely closed economy,

(a) PPP holds relative to the surrounding countries.

(b) A 10 percent devaluation of the host currency will be offset by a 10 percent rise in the host country prices.

(c) The value of a foreign subsidiary, in units of the foreign parent's home currency, is unaffected by exchange rate changes.

(d) The real value of a foreign subsidiary to an investor from the host country is unaffected by exchange rate changes.

(e) In the absence of contracts with a value fixed in host currency, the real value of a foreign subsidiary to an investor from the parent's home country is unaffected by exchange rate changes.

(f) In the absence of contracts with a value that is fixed in foreign currency, the real value of a foreign subsidiary to an investor from the host country is unaffected by exchange rate changes.

(g) There is little or no advantage to having one's own currency: exchange rate policy has virtually no effects.

3. In an economy that is neither perfectly open nor completely closed,

(a) Consider a company that produces and sells in this economy. Apart from contractual exposure effects, its value in terms of its own (local) currency is positively exposed to the value of other currencies.

(b) The value of an importing firm located in this economy could either go up or go down when the local currency devalues: the effect depends on such factors as the elasticity of local demand and foreign supply.

(c) Consider a company that produces and sells in this economy. Apart from contractual exposure effects, its value in terms of a foreign currency is positively exposed to the value of its currency expressed in terms of other currencies.

4. Suppose that the value of the firm, expressed in terms of the owner's currency, is a linear function of the exchange rate up to random noise.

   (a) The firm's exposure is the constant $a_{t,T}$ in $V_T(i) = a_{t,T} + b_{t,T}\, S_T(i) + e_{t,T}(i)$.
   (b) The exposure is hedged by buying forward $b_{t,T}$ units of foreign currency.
   (c) Hedging means that all risk is eliminated.

5. Suppose that the value of the firm, expressed in terms of the owner's currency, is a nonlinear function of the exchange rate up to random noise. Suppose that you fit a linear regression through this relationship, and you hedge with a forward sale with size equal to the regression coefficient.

   (a) All risk will be eliminated.
   (b) There is remaining risk, but it is entirely independent of the realized value of the exchange rate.
   (c) There is remaining risk, but it is uncorrelated to the realized value of the exchange rate.
   (d) There is no way to further reduce the variance of the firm's hedged value.
   (e) There is no way to further reduce the variance of the firm's hedged value if only exchange rate hedges can be used.
   (f) There is no way to further reduce the variance of the firm's hedged value if only linear exchange rate hedges can be used.

13.6.2 Applications

SynClear, of Seattle, Washington, produces equipment to clean polluted waters. It has a subsidiary in Canada that imports and markets its parent's products. The value of this subsidiary, in terms of CAD, has recently decreased to CAD 5m due to the depreciation of the CAD relative to the USD (from the traditional level of USD/CAD 0.85 to about 0.75). SynClear's analysts argue that the value of the CAD may very well return to its former level if, as seems reasonable, the uncertainty created by Canada's rising government deficit and Quebec's possible secession is resolved. If the CAD recovers, SynClear's products would be less expensive in terms of CAD, and the CAD value of the subsidiary would rise to about 6.5m.

1. From the parent's (USD) perspective, is the exposure of SynClear Canada to the USD/CAD exchange rate positive or negative? Explain the sign of the exposure.

2. Determine the exposure, and verify that the corresponding forward hedge eliminates this exposure. Use a forward rate of USD/CAD 0.80, and USD/CAD 0.75 and 0.85 as the possible future spot rates.

3. SynClear's chairman argues that, as the exposure is positive and the only possible exchange rate change is an appreciation of the CAD, the only possible change is an increase in the value of the subsidiary. Therefore, he continues, the firm should not hedge: why give away the chance of gain? How do you evaluate this argument?

   In the remainder of this series of exercises, SynClear Canada's cash flows and market values are assumed, more realistically, to depend on other factors than just the exchange rate. The Canadian economy can be in a recession, or booming, or somewhere in between, and the state of the economy is a second determinant of the demand for SynClear's products. The table below summarizes the value of the firm in each state and the joint probability of each state:

State of the economy	Boom	Medium	Recession
S_T=0.85:Joint probability	0.075	0.175	0.25
Value_T (USD)	5.25	4.75	4.50
S_T=0.75:Joint probability	0.25	0.175	0.075
Value_T (USD)	4.25	3.857	3.50

4. What are the expected cash flows conditional on each value of the exchange rate?

5. Compute the exposure, the optimal forward hedge, and the value of the hedged firm in each state. The forward rate is USD/CAD 0.80.

Part IV

Long-Term International Funding and Direct Investment

About this Part

The prime sources for long-term financing are the markets for fixed-interest instruments (bank loans, bonds) and stocks. We review the international aspects of these in Chapters 16 and 17-18, respectively. Expected returns on stocks provide one key input of investment analysis, so in Chapter 19 we consider the CAPM and the adjustments to be made to take into account real exchange risk. The other inputs into NPV computations are expected cashflows, and these are typically quite similar to what one would see in domestic projects. There is one special issue here, international taxes (Chapter 20). In Chapter 21 we see how to do the actual NPV, extending the usual two-step approach—NPV followed by Adjusted NPV to take into account the aspects of financing, relevant in imperfect markets—to a three-step version to separately handle intra- and extra-company financial arrangements. We conclude with an analysis of joint-venture projects, where NPV is mixed with the issue of designing a fair profit-sharing contract (Chapter 22).

Case: A JV Project Between Weltek, Antwerp and Fusioneering, Jamspedpur

This case text is based on a real-world case, but the names of the two main companies and their managers, as well as all dates and amounts, have all been changed.

Mr. Dondeyn is the General Manager of Weltek, a producer of welding electrodes and equipment. Over the last years, he and his Assistant General Manager Ms Dewulf have been negotiating a joint-venture with three Indian partners, aiming at the local production of electrodes and possibly also the distribution of Weltek products imported into India.

The core business of Weltek is in “special” welding electrodes for maintenance and repair, not for plain-vanilla construction welding. Weltek belongs to the subtop in the industry and would like to grow. Founded in Belgium in the 1960s, it has subsidiaries in Italy, the Netherlands, the UK, Spain, the US, and South Africa. All these are wholly-owned. Production is concentrated in Belgium and Spain; the other subsidiaries are marketing and service companies.

Weltek has been interested in an Indian production unit for years. The internal market is huge, not only because of the size of the population, but more importantly also because repair is big business there. Like many developing countries, India is short of capital to import new equipment, and in the late 1990s the still-highish import tariffs make replacement very expensive; as a result, most equipment (industrial machinery, cars, appliances, etc.) is used much longer than in most OECD countries. This implies an important maintenance and repair market, which in its turn induces a market for hand-welding fillers and equipment. India could also be a stepping-stone for exports towards other countries in the area, including CIS countries.

Weltek had in mind a production joint venture, not a wholly owned subsidiary nor a marketing joint venture. The JV option had to do not only with the local regulations—the investment code limited foreign ownership to 40%—but also with Weltek’s financial capacity and its lack of knowledge of the local market. There was a transfer-risk issue too: in those days, there was free repatriation of capital brought in for direct investments, but there still was a bureaucratic delay, and the occasional nationalist noises by the then ruling BJP party were not encouraging. (Restrictions on portfolio investments were even more stringent, whether in- or outward, in those days.) A JV meant a smaller investment, so smaller transfer risk. Local production, not just marketing of imported products, was preferred because even after India’s liberalization of the 1990s its import tariffs are still high by Western standards. Therefore, exports to India are viable only for selected specialty products for which demand is too low to justify local production. A pure license contract wouldn’t do either. One argument was control over the training and marketing effort. “We have a very intensive and well-developed formal training scheme for the sales force and for the engineers; an engineer, for instance, has to know literally everything about his or her products. So we were not willing to surrender control over training to an independent party", Mr. Dondeyn explains. Second, in view of the Government restrictions on the life of a license contract (five years) and on the size of the payments (at most 5%), a joint venture could be expected to be much more profitable than a stand-alone license contract.

Fusioneering, of Jamshedpur, was, amongst others things, in the arc-welding business, but given the youth of the company and its as yet limited expertise its managers felt that they needed more up-to-date know how and technology. Thinking that the really big fish would not be interested in a small and as yet unprofitable business, they talked to mid-sized players like Weltek. A visit led to lengthy policy discussions, and ultimately to a Memorandum of Understanding. The parties agreed to live with low or negative profits for a couple of years, during which free samples would be distributed to build up credibility. The JV would buy the raw materials for production of ten Lastek electrodes, and sell the output. Still, it took almost two more years until all feasibility checks had been done and (especially) all investment licenses were obtained. During that time, Mr. Dondeyn was busy with other things. The Asian financial crisis and a misunderstanding about registration put the whole project into the freezer, but talks were reopened after a chance meeting at a trade fair. The tentative deal was as follows.

1. Investments are estimated and timed as follows:
    - Land is bought and paid for on 1/1/2000; cost: 5m rupees
    - Construction (Plant and Equipment) starts on 1/4/2000 and lasts six months; cost: 10m
    - Training of engineers starts on 1/7/2000; cost: 5m including travel to and from Belgium

   Total upfront investment is, therefore, 20m. Also to be financed are the initial cash drains, estimated at about 5m in year 0 and 3m in year 1 (see below).

   The equity is set at 10m, 40% of which is provided by Weltek and 60% by Fusioneering. The equity is fully paid up on 1/1/2001 but lent back to the shareholders at zero cost with the proviso that they should finance, 40%/60%, any cash drain up to, cumulatively, 10m that would occur; if cash drains exceed 10m, any shareholder has the option to sell his stake at book value to the other; and if both want to sell, the JV is to be liquidated. The investments themselves are financed by a loan (20m rupees at 8%), guaranteed entirely by Weltek and Welteks house bank; the loan is to be taken up on 1/1/2000, and amortized in five equal annual tranches of 4m as of (the end of) 2001.

2. Production starts 1/10/2000. There is a 3-month lead time (production, storage in Hyderabad, distribution, storage at the sales point) between the start of production and sales, so sales start on 1/1/2001. The contract also stipulates that Weltek receives an up-front license fee of 2m rupees, plus a five-year, 5% royalty on net sales payable early April after the reporting period.

Table 15.3: **Pro Forma P&L statements**

time (year)	0	1	2	3	4	5
Investments						
Land	5					
Plant and Equipment	10					
Training	5					
Upfront License	2					
P&L Projections (1999 prices)						
Revenue:						
Gross Sales	0	10	17	22	24	25
<excise tax>	0	0.3	0.51	0.66	0.72	0.75
<rebates>	0	0.5	0.85	1.1	1.2	1.25
<provisions bad debt>	0	0.15	0.255	0.33	0.36	0.375
Net Sales	0	9.05	15.385	19.91	21.72	22.625
Cost:						
Variable Production Costs	0	4	6	7	8	8.5
Depreciation P&E	2.5	2.5	2.5	2.5		
Depreciation Training	2.5	2.5				
Overhead	1	2	2.5	3	3	3
License: Upfront Fee	2					
Royalties	0	0.46	0.782	1.012	1.104	1.15
Interest Paid	1.6	1.6	1.28	0.96	0.64	0.32
Profit Before Tax	-9.6	-4.01	2.323	5.438	8.976	9.655
Tax computations						
Cumulative profits	-9.6	-13.61	-11.287	-5.849	3.127	12.782
Losses carry Back	-9.6	-13.61	-11.287	-5.849	0	0
Taxable income	0	0	0	0	3.127	9.655
Taxes	0	0	0	0	1.2508	3.862
Profit After Tax	-9.6	-4.01	2.323	5.438	7.7252	5.793
Cash flows to shareholders						
Add Back Depreciation	5	5	2.5	2.5	0	0
<Amortization Of Loan>	0	4	4	4	4	4
Cash Flow	-4.6	-3.01	0.823	3.938	3.7252	1.793

3. Projected P/L are based on the following figures listed below and the following addenda:
   - Sales in the first year start as of 1/1/2001. Gross sales is computed as (volumes sold) × (list prices). The list prices include the excise tax (3%), which must be deducted for the purpose of profit calculations. Also deducted is an estimated 5% representing rebates for large orders, and a provision for non-performing receivables estimated at 1.5%. The result is net sales income. Customers obtain a 30-day credit period.
   - Production costs (variable, depreciation, overhead) are as shown in the table. Training can be depreciated over two years, and equipment over four years, starting in year zero. Depreciation has to be linear.
   - Know-how and financial charges. There is a 2m upfront licensing fee and

the 5% royalty on net sales. The interest on the bank loan is computed on a loan balance of 20m in January 2000 and 2001, 16m in January 2002, 12m in January 2003, 8m in January 2004, and 4m in January 2005. Interest is payable in four quarterly tranches (2% effective per quarter).

- Taxable profit is the annual profit minus any tax shield from carried-over losses. Taxes (40%) are payable in the middle of the year following the reporting year; that is, taxes on 1994 profits are paid mid-1995 etc.

On the flight back to Belgium, Mr Dondeyn types the projected P/L statement into his laptop, and runs a quick-and-dirty NPV. To compute the cash flows, Mr Dondeyn notes that the net investment is zero (total investments are entirely financed by a loan); so he takes profits after taxes, adds back depreciation and subtracts the loan amortizations. The cost of capital is set at 14% (the 8% on the loan being the subsidized rate at which the JV would be able to borrow, plus a 6% risk premium assuming a unit beta). The NPV seems to be –0.71m rupee—not hugely negative, but negative nevertheless.

Issues

1. The NPV calculations do not seem to involve anything special: it all looks like domestic capital budgeting. Are there no special issues that arise when the project is international?

2. The quick-and-dirty calculation ascribed to Mr Dondeyn ignores the fact that there are two shareholders, and is made as if this were a wholly-owned project. Even so, these calculations are very flawed: I stuffed about every mistake into the spreadsheet one could possibly make. Read this part, identify the errors, and correct them.

3. Is there a way to judge the fairness of the proposed cash-splitting rules for the JV? What should one look at? If one finds a good measure of fairness, is there a straightforward way of achieving this fairness, or is it just a matter of trial and error?

Chapter 16

International Fixed-Income Markets

In this chapter, we have a look at one source of financing for companies: international money, loan, and bond markets. Related short-term fixed-interest products, like deposits and short-term loans or commercial paper, are briefly touched upon. Other related instruments, notably interest forwards and futures, were introduced in the Appendices to Chapters 4 and 6.

These international fixed-interest markets used to largely coincide with what was (and largely still is) called *euromoney* and *eurobond* markets, that is, markets for banking products or bonds denominated in a currency that is not the official money of the country where the loan was taken up or the bonds were issued. For example, a Norwegian investor may deposit USD not in the US but with a bank located outside the US, for example, in Oslo or in London. Or a Peruvian company may issue bonds in London and denominate them in JPY.

The prefix "euro" became misleading when such extra-territorial markets also emerged in e.g. Asia. One accordingly heard of Asiadollars, and so on. Since the advent of the euro as a currency, the prefix has also become ambiguous: are we talking about bonds expressed in EUR or bonds issued outside the home turf of the currency? Also, the term could lead to absurd combinations, like euro-euro for EUR bonds placed in London. There have been feeble attempts to find a new term; *The Economist* even invited suggestions from the public at large, and in the end backed the by no means new "xeno" proposal (from greek $\xi\epsilon\nu o\varsigma$, foreign). But the entire prefix issue fizzled out on its own, since people no longer thought there was anything special about setting up deals in a particular currency outside its original territory. If ever the distinction is important to you, you can just add the adjective: everybody will catch your drift. In most of this text I prefer to use "international". The term "offshore" might have done well, too, if it weren't for the connotation with "having a special tax status", which is not what we have in mind right now.

In the sections that tell the tale of how the international markets emerged, we still use the euro- prefix in its "international" meaning. For simplicity, when we say euro- we also mean to include other international markets in the Middle East and especially in the Far East (Tokyo, Hong Kong, Singapore). This largely conforms to standard practice in the Americas and Europe.

The earliest activity in the international markets was in the deposit and loan segments, the segments where banks act as intermediaries between investors and borrowers. The emergence and growth of the eurobanking business was mainly the result of low costs, which enable a more narrow bid-ask spread. The success of this unregulated, wholesale banking market was soon imitated in the bond section and in the short-term securities part of the capital market (eurobonds and eurocommercial paper, respectively).

This chapter is organized in the following way: In the first section, we describe the traditional eurobanking world: markets for short-term international deposits, bank loans, and credit lines. We then discuss the counterparts of these banking products in the securities markets, namely, the international bond and commercial paper (CP) markets (Section 2). The final issue we bring up, in Section 3, is how to compare one's fixed-financing alternatives across currencies and markets. We conclude in Section 3.

16.1 "Euro" Deposits and Loans

The banking segment is the oldest segment of the international markets. Even before World War II, there was a small market for USD deposits and loans in London, then the world's financial heart. However, the market took off in earnest in the sixties only. We start by explaining the reasons for its rapid growth since then. We distinguish between circumstances that facilitated the emergence of the market and reasons for its longer-term success. The proximate reasons had mostly to do with bad economic-policy decisions and regulations that had unexpected consequences.

16.1.1 Historic, Proximate Causes of Euromoney's growth

Liberalization of trade and exchange The eurodollar markets began to expand in the fifties and sixties, after the lifting of the widespread exchange controls. These controls had been imposed after World War II because of the scarcity of dollars (the only internationally accepted currency at the time, since even the GBP's international use had become heavily controlled and regulated).

Note, however, that while liberalization of the exchange market is a necessary condition for the emergence of euromoney markets, it is not an explanation of that emergence.

The US trade deficit The liberalization of the European exchange markets was possible only because the shortage of USD did not last long. Immediately after the war, the US launched an international aid program (the Marshall Plan). In addition, the US imported more goods and services than it exported, and US corporations became important international investors, buying companies or building plants all over the world. But, tautologically, the balance of payments has to balance, remember. So the deficit on the current account, the "capital" (aid) account, and the FDI balance meant that there must be a surplus or a set of "source" deals, elsewhere (see Chapter 2). This offsetting surplus was realized by exporting US government or corporate bonds and short-term assets, including sight money. Most countries cannot export sight money in any meaningful amounts, but the US could since its money was also the closest one can get to world money: it was used everywhere for international transactions. Thus, the US' deficit on the current account and the aid and FDI accounts meant that more and more sight money ended up in the hands of foreign investors. Foreign central banks held some too, but prefered interest-bearing forms.

Note, again, that this is not a true explanation for the rise of euromoney markets. The fact that there were foreign-owned USD does not explain why part of these USD balances were held via European banks rather than directly with US banks. The next three arguments relate to positive incentives for eurodollar transactions.

Political risks Since the fifties, the cold war created political risks for communist countries that wished to hold USD deposits: the US government could seize Soviet deposits held in New York. For that reason, the Soviet Union and China shifted their dollar balances to London and Paris, out of reach of the US government. This meant that there was a Western bank between them and the US banking system (see, again, Chapter 2).

UK capital controls and restrictions In the nineteenth century, London had been the world's center for international financing and Sterling the world's favorite currency. After World War II, however, the GBP was chronically overvalued, and the UK had serious balance-of-payments problems of its own.[1] Thus, the British government limited foreign borrowing in GBP.[2] As a result, UK banks borrowed USD

[1]Until 1949, the GBP maintained the gold parity that was originally fixed in 1752 by Isaac Newton. (Newton was director of the Royal Mint, a sinecure job meant to leave him time for research.) This had become a very unrealistic rate in view of the increase in paper-money supply after World War I and, especially, World War II. So British exporters had a hard time while imports boomed. The UK could no longer hope that sterling balances, sent abroad in payment for its net imports, would happily be held by traders all over the world; rather, outstanding pound balances were being returned and converted into dollars. All this put pressure on Sterling. (The pound devalued by 40 percent in 1949, and by another 16 percent in 1967. Controls were lifted gradually, and entirely went out of the window only in the 1980s, under Thatcher.)

[2]A speculator, remember, would borrow a currency deemed weak, and sell the proceeds (hoping to be able to buy back later at a low cost), thus putting additional pressure on the spot rate. Until the 1980-90s, governments often had the lamentable habit of forbidding the symptons rather than curing the disease. So HM's government forbade pound loans to non-residents.

(that is, accepted USD deposits), which were then used to extend USD loans instead of GBP loans.

US capital controls and restrictions In the US, the disequilibrium on the merchandise&invisibles, aid, and direct-investment balances, combined with the growing overvaluation of the USD against the DEM and related currencies, pushed up US interest rates. President Kennedy tried to alleviate the problem by imposing the *Interest Equalization Tax* (1963–74) on foreign borrowing in the US. This tax allowed internal US interest rates to remain below USD interest rates offered in Europe. President Kennedy and later, President Johnson, also imposed *foreign credit restraints* (1965, 1968–74) that hindered borrowing by foreigners in the US market. Simultaneously, *Regulation Q* (enacted in 1966, relaxed in 1974, and abolished in 1986) imposed interest ceilings on domestic USD deposits with US commercial banks. The combined effect of all of this was that US corporations and investors preferred to hold USD deposits in Europe (where they obtained better rates), and these dollars were then re-lent to non-US borrowers who were no longer allowed to borrow USD in the US. Finally, President Nixon's "voluntary" (and later, mandatory) *curbs on capital exports* had the unintended result that US multinationals avoided depositing their funds in the US lest these funds be blocked there. The money were deposited, instead, in the euromarkets.

16.1.2 Comparative Advantages in the Medium Run

The eurodollar markets did not collapse after all of the regulations described above were abolished. Nor can the above factors explain the subsequent emergence of international markets for other currencies, like the DEM, the JPY, or the ECU, and—to a lesser extent—the CHF, NZD, NLG, etc. The long-term explanation of the success of these international markets is their lower bid-offer spread (that is, the difference between interest rates on loans and interest rates on deposits), which in turn reflects the lower costs of international banking as compared to domestic banking. There are many reasons for the low operating costs:

A lean and mean machine The international market is essentially a wholesale market, where large volumes of transactions allow narrow spreads. Eurobanks, unlike many domestic commercial banks a few decades ago, were not expected to offer politically or socially inspired subsidized loans to ailing companies or house-building families. Nor do they need an expensive retail network.

Low legal costs Most euroborrowers are sovereign states or high-grade corporations. This means that there are hardly any costs of credit evaluation, bonding, and monitoring.

Lighter regulation For eurodollar banking (as opposed to domestic banking) there is no compulsory deposit insurance, which means that there are no insurance costs. Nor are there any reserve requirements (which are, in fact, similar to taxes on deposits),[3] and local monetary authorities tended to be far more lenient as far as credit restraints are concerned when borrowing did not involve their home currency.[4]

Universal banking In the UK, like in much of continental Europe, there was no equivalent of the US Glass Steagall Act that separated commercial banking (sight and time deposits; overdrafts and loans) from investment banking (placing, underwriting, and holding securities). Although, by definition, you do not need universal banks for deposits and loans, companies still liked institutions that could both offer loans and help place their paper: both are very close substitutes. Nor was there anything, in the UK, like the US' ban on interstate banking, a rule that imposed a cap on US banks' growth (except for a handful of international players).

Lower Taxes Eurobanks were often located in tax havens or are part of a financial network involving tax havens. Also in mainstream OECD countries, international transactions often received beneficial tax treatment when compared with domestic businesses (for example, a waiver of stamp duties or withholding taxes; in this respect, many OECD countries have now followed the lead of tax-haven countries).[5] Furthermore, many investors with undeclared income—the "Swiss dentist" or the "Belgian dentist", as *The Economist* or *Euromoney* fondly call them—appreciated the opportunities for tax avoidance or tax evasion available in euromoney markets. Foreign deposits were often fiscally anonymous (that is, the bank cannot be forced to reveal their identity to a foreign tax authority), or are often in the form of bearer securities.

16.1.3 Where we are now: a Truly International Market

As Merton Miller beautifully put it, silly regulation provided the grain of sand—the thing that starts as an irritant to an oyster but ultimately grows into a pearl.[6] The market survived the abolishment of the currency controls and excessive regulation; these had speeded up the growth of the market, but even without them a

[3] A 5 percent reserve requirement would mean that a bank, when receiving a customer deposit for 100, has to redeposit 5 in a non-interest-earning account with the central bank. Thus, only 95 can be re-lent. This increases the effective cost of accepting the deposit.

[4] It is true that eurobanks are subject, like any other banks, to the so-called Bank of International Settlements (BIS) rules. However, these are not reserve requirements. Rather, the BIS rules set the minimum amount of equity a bank should have, in light of its balance-sheet and off-balance-sheet positions.

[5] A stamp duty is a tax on transactions in securities. A withholding tax is a tax levied on interest or dividends, withheld when the interest or dividend is paid out (rather than collected afterwards, on the basis of a tax return).

[6] Actually, pearls do not grow in oysters (*ostreida*) but in two species closer to mussels, *pterida* and *meleagrina* (aka *pintada*); and they do not start from grains of sand but from indigestible food residues. But otherwise the image is accurate.

Table 16.1: **Largest Banking Centers**

	Banking deposits Dec-05 (USD tr)	# of banks Mar-05 (number)	Branches & subsidrs of forgn banks Mar-05 (number)	Return on captl 2006 (%)	Cost/ income ratio 2006 (%)	Cross-border lending (%) Mar-07	Cross-border borrowing Mar-07 (%)
JP	5.1	129	69	16.2	69.4	9	2
US	5.2	7,540	228	22.2	59.9	20	12
UK	4.6	347	264	19.6	50.2	10	23
DE	3.1	**2,344	125	4.7	61.5	9	7
FR	1.5	**318	*217	15.7	63.9	8	9

* end-2004; ** end-2005

Source International Financial Services London. *International Financial Markets in the* UK , November 2007. www.ifsl.org.uk/research; based on Bank of England, KPMG, Banque de France, The Banker, Bank for International Settlements, European Banking Federation, Bank of Japan, FDIC, US Federal Reserve

similar market would have emerged sooner or later because there was a need for a fast, lightly regulated field for big, professional players. London and other centers provided just that.

Nowadays, however, the playing field has become much more even, and the “euro”markets’ comparative advantage was eroding fast. Wholesale, simple deals with prime borrowers can be signed everywhere. “Regulatory arbitrage”—that is, borrowers and investors migrating away from the overly regulated markets—has forced countries everywhere to dump rules, taxes and duties that did more harm than good. In the US, Glass-Steagall and the ban on interstate banking have been repealed. Currency controls are gone for most currencies, as are credit restraints and, in many countries, reserve requirements. Tax authorities cooperate internationally, and governments exchange information on foreign deposits and/or foreign investment income. Originally, this was just in cases where crime- and drugs-related money laundering was suspected or, later, terrorist activities; but cooperation for fiscal purposes is coming in too. Within the EU, information on non-residents’ income is already being shared; a few countries still impose withholding taxes instead but this will be phased out. Secret (“numbered”) bank accounts, for a long time one of the attractions of, most notably, Swiss, Austrian, and Lichtenstein banks, are essentially a thing of the past: bankers must know their customers’ identities. Tellingly, countries with a dark reputation are now being blacklisted by the OECD; when in the early 2000’s a new government in Mauritius proposed to set up a “high-privacy” banking sector, the big countries were all over her and Mauritius hastily withdrew the proposal.

As a result, there is not much difference anymore between domestic and international banking, certainly not for wholesale deals in OECD countries and the like. In that sense, in a large part of the world markets nowadays are truly international, not just a collection of local markets with, on the fringe, an international corner for the big guys.

In the following sections, we review the products offered by international banks. The first product in our *tour d'horizon* is the deposit.

16.1.4 International Deposits

Initially, international deposits were typically *time deposits* (or *term deposits*)—that is, non-negotiable, registered instruments with a fixed life. A *certificate of deposit* (CD) is the tradable-security version of the traditional term deposit: it is negotiable (that is, can be sold to another investor at any time) and is often a bearer security.

The bulk of the deposits have a very short duration—for instance, overnight, one or two weeks, but mostly one, three, or six months. These short-term deposits or CDs pay no interim interest; there is a single payment, principal and interest, at expiration. For long-term CDs or long-term deposits (up to seven years), there is a fixed coupon or floating-rate coupon. For CDs with *floating-rate* coupons, the life of the CD is subdivided into subperiods of usually six months. The interest rate that applies for each period consists of a fixed spread laid down in the contract, and a risk-free market rate that is reset every period. Following the by now familiar "spot" tradition, this re-setting occurs two working days before the beginning of the period (the *reset date*). The market rate on the basis of which the rate is reset is usually the *London Interbank Offer Rate (LIBOR)* or the Interbank Offer Rate in the currency's domestic financial center. "The" LIBOR and similar -IBOR's[7] are computed as an average of the rates offered by an agreed-upon list of banks; the EBA has standard lists. The basis of the floating rate may also be the bid rate, or the mean (midpoint) rate, or, in the US, the T-bill rate or the prime rate. If the basis rate is an ask rate (like IBOR or the prime rate[8]), the spread is usually negative: we are talking about deposits, here.

Example 16.1

An investor buys a NZD 1,000,000 floating-rate CD with a life of two years, at NZD LIBOR minus 0.375 percent, reset every six months. The initial reference interest rate is 4 percent *p.a.*, which implies that after six months the investor receives $1,000,000 \times (4 - 0.375)\%/2 = \text{NZD}18,125$. The reset date is two days before this interest is paid out, and the six-month LIBOR on this reset date may turn out to be, say, 3.5 percent *p.a.* This means that the second interest payment will be only

[7] The rule is to make the first letter refer to a city, like Aibor (A'dam), Bibor (Brussels), Pibor (Paris), Tibor (Tokyo) and so on. For the EUR one refers to Euribor; the old national IBORs have gone, including Frankfurt's Fibor.

[8] The prime rate once was the posted rate for unsecured loans to good-quality borrowers; nowadays it is, de facto, applied to rather mediocre borrowers.

$1,000,000 \times (3.5 - 0.375)\%/2 =$ NZD 15.625. There will be two more of these reset dates. At the end of the last period, the principal is also paid back.

You can view such a floating-rate CD as a series of short-term CDs that are automatically rolled over without reinvestment of the interest earned each period. Sometimes a floating-rate CD has a *cap* or a *floor*—that is, the interest rate that the investor actually receives has an upper or lower bound. We shall discuss caps and floors in the next section, which describes euroloans.

16.1.5 International Credits and Loans

International banks offer essentially the same products as domestic banks: loans and credit lines. But there are a few interesting differences.

Consortia

One difference is that the loans tend to be extended by a group of banks (a *syndicate* or *consortium*) rather than by a single institution. The members of the consortium or syndicate can play very different roles:

- The *mandated arranger* (or, more traditionally, the *lead bank* or *lead manager*) negotiates with the borrower for tentative terms and conditions, obtains a mandate from him to get the loan together, and looks for other banks that are willing to provide the money or at least to act as a back-up for the money. In the event that a group of banks unanimously agree to form a mandated arranger group, the title is assigned to all banks within such group. *Book runner* status is then assigned to the bank that runs the book (i.e., sollicits and records the commitments by other banks to participate in the funding) for a deal that is sent out into general syndication. The book runner often leads the consortium even if arranging is, formally, shared. Bookrunnership is now, in turn, getting shared among many banks; soon we will need a lead bookrunner, and a few years from now a coordinator of the lead bookrunners.
- The banks that provide the actual funding are called *participating banks*.
- Because the funding is not yet arranged at the time of the negotiations, the lead bank or the group of joint bookrunners often contacts a smaller number of *managing banks* to underwrite the loan, that is, guarantee to make up for the shortage of funds if there is a shortage. These banks are also called *underwriters* or *co-managers* or *co-leads* or *co-arrangers*.
- The *paying agent* or *facility agent*, finally, is the bank that receives the service payments from the borrower and distributes them to the participating banks.

Any given bank can play multiple roles. For instance, the lead bank is almost invariably also the largest underwriter (hence, the name “lead manager”) and usually provides some of the funding as well. The main objective of syndication is to spread the risks, but it also eliminates the moral hazard of the borrower paying off its bigger lenders and ignoring the small debt holders: because of the paying agent system, the borrower either defaults toward all banks, or toward none.

As in domestic banking, the borrower often signs promissory notes (that is, I owe you [IOU] documents), one for each payment. The advantage of receiving promissory notes is that they are easily negotiable. That is, if the lending bank needs funds, it can pass on the promissory note to another financial institution as security for a new loan, or it can sell the promissory note. Regular loans are not so easily traded: they need to be packed into special vehicles which then issue claims against the vehicle's assets (loan-backed securities, collateralized debt obligations, and the like).

Until well into the 1990s a big loan would show up in *Euromoney*, *The Economist*, or *Business Week* and the like as a "tombstone"—that is, an austere-looking advertisement that trumpets the signing of a new deal. Sadly, these are now replaced by internet press releases. Figure 16.1 shows one by Turkey's Vakifbank.

Revolving or Floating-Rate Loans

Another difference between traditional bank loans and big international loans is that the latter tend to be of the floating-rate (FR) type, whereas many domestic loans still have an interest rate that is fixed over the entire life of the loan. The reason for this predilection for FR loans is the very short funding of banks (see above): banks do not like the risk that, after having lent long-term at a fixed rate, they may have to refinance short-term at unexpectedly high interest rates. The emergence of interest swaps, however, has made the hedging of an interest gap easier. International banks now lend longer and domestic banks resort to FR loans more often too.

Example 16.2

A bank accepts a three-month, DKK 100m deposit at 4 percent *p.a.* and extends a loan for six months at 4.5 percent *p.a.* For simplicity, assume that this deposit and this loan represent the bank's entire balance sheet. After the deposit has expired, the bank must pay DKK 100m $\times$ (1 + 4%/4) = DKK 101m to the original lender. Since there are no cash inflows yet from the loan, the bank must borrow this amount (that is, accept a new three-month deposit). If, at that time, the three-month rate has increased to 7 percent *p.a.*, then after another three months the bank has to pay 101 $\times$ (1 + 7%/4) = DKK 102.767,5m though it receives only 100m $\times$ (1 + 4.5%/2) = DKK 102.250m from the original six-month borrower. Thus, because of the increase in the short-term interest rate the bank lost DKK 517,500 rather than making money.

In the above example, the maturity mismatch is not large because the loan is assumed to be for only six months. However, borrowers often have long-term capital needs; and rolling over short-term loans (at interest rates revised at each roll-over

Figure 16.1: **An international syndicated loan: Vakifbank**

$500m Turkish bank loan syndication for Vakifbank signed in Dubai

UAE Central Bank governor addresses signing ceremony

6 December 2004; Dubai, UAE: A US$500 million syndicated term loan agreement for Vakifbank (Türkiye Vakiflar Bankasi T.A.O.) one of the strongest banks in Turkey, was signed in Dubai today by a syndicate of 56 blue-chip regional and international banks. The loan was raised to pre-finance Turkish export contracts and has a margin of 60 basis points per annum. VakifBank is currently rated Bpi by S&P, B+ by Fitch, and B2 by Moody's. On 1st of November 2004 Fitch increased the National Long Term Rating of VakifBank by two notches to A-(tur).

Bookrunners: Citibank NA, Standard Bank London Limited and WestLB AG

Documentation Agent: Standard Bank London Limited

Facility Agent: Sumitomo Mitsui Banking Corporation Europe Limited

Information Memorandum: WestLB AG

Coordination, Publicity and Signing: Standard Bank London Limited

Mandated Arrangers: ABN AMRO Bank N.V., Al Ahli Bank of Kuwait, Alpha Bank A.E., American Express Bank GmbH, Banque Saudi Fransi, The Bank of Tokyo-Mitsubishi, Ltd., Burgan Bank, Citibank N.A., Demir-Halk Bank (Nederland) N.V., Deutsche Bank AG London; Dresdner Kleinwort Wasserstein (acting through Dresdner Bank AG, Niederlassung Luxemburg), GarantiBank International N.V., Gulf Bank KSC, HVB Group (represented by members of HVB Group), ING, J.P. Morgan plc, Mashreqbank P.S.C., Natexis Banques Populaires, Raiffeisen Zentralbank Österreich AG, Standard Bank London Limited, Standard Chartered Bank, Sumitomo Mitsui Banking Corporation Europe Limited, UFJ Bank Limited; Wachovia Bank, National Association and WestLB AG, London Branch.

Co Arrangers: Gulf International Bank B.S.C., HSH Nordbank AG, Samba Financial Group, , Managers, Doha Bank, Arab African International Bank, Banque Misr – Overseas Branch, Erste Bank (Malta) Limited, Finansbank (Holland) N.V., Raiffeisenlandesbank Oberösterreich Aktiengesellschaft, The Bank of Nova Scotia, The Commercial Bank Of Qatar (Q.S.C.), The Saudi National Commercial Bank, Bahrain, UBAE Arab Italian Bank Spa

Participants: Bankmuscat S.A.O.G., Arab Bank plc, Baden-Württembergische Bank Aktiengesellschaft, Banca Nazionale del Lavoro S.p.A., London Branch, Banco Bilbao Vizcaya Argentaria S.A., Bank Hapoalim B.M., Bank of Ireland, Banque Internationale De Commerce – BRED, Credit Suisse, First Gulf Bank, Misr International Bank S.A.E., Sabanci Bank PLC, United Bank Limited, UAE, Z_ivnostenská Banka, a.s., Banca Monte Dei Paschi di Siena s.p.a., Habib Bank AG Zürich, London Forfaiting Company Limited, Tunis International Bank

Source http://www.international.standardbank.com/AboutUs/PressReleases.aspx?id=232

date) is awkward for the borrower: the bank can always refuse to extend a new loan or drastically increase the spread over LIBOR. The need to reconcile the banker's desire for a safe interest margin with the borrower's preference for long-term guaranteed funding gave rise to the *revolving loan* or *floating-rate loan.* This is a medium-term or long-term loan where the interest rate is reset every period on the basis of the then-prevailing money market rate plus a spread. For example, if interest is payable every six months, then, just like for FR deposits, two days prior to the beginning of each such period, the interest rate for the next half-year is set equal to the then-prevailing six-month LIBOR rate, contractually increased with a spread of, say, 1/2 percent *p.a.* Thus, the bank is protected against interest risk, and the borrower's funding is guaranteed for an agreed-upon period at a preset risk-spread over the base rate. The basis of the interest rate in rolled-over loans is typically the LIBOR or a similar interbank rate. Occasionally, the US prime rate or the US T-bill rate is chosen.

Table 16.2: **A revolving loan with a cap and a floor**

rates (%)		equivalent PN story			mood		
LIBOR	rate on loan	face value	fair PV	proceeds	bank	you	option ex'ed
3	$3.5/2 = 1.75$	101.75	100.25	100	☺	☹	bank's call
3.5	$3.5/2 = 1.75$	101.75	100.00	100			(none)
4	$4.0/2 = 2.00$	102.00	100.00	100			(none)
4.5	$4.5/2 = 2.25$	102.25	100.00	100			(none)
5	$4.5/2 = 2.25$	102.25	99.76	100	☹	☺	your put

Key The loan is about 100m, at 6-month LIBOR but with a floor at 3.5% *p.a.* and a cap at 4.5% *p.a.*. Thus the six-monthly PN the borrower has to hand over has a face value of 100×(1+LIBOR), but no lower than 101.75m and no higher than 102.25. With very low LIBOR s, when the fair value of a PN at 101.75 would be higher than 100, the borrower gets no more than 100: the bank exercises its call on the Note. With very high LIBOR s, when the fair value of a PN at 102.26 would be lower than 100, the borrower still gets 100: she exercises her put on the Note.

Revolving Loans with Caps or Floors

Sometimes there is a *cap* and/or a *floor* to the effective interest rate. For instance, the contract may say that the interest rate will never exceed 6 percent *p.a.* (cap), or fall below 3 percent *p.a.* (floor). These caps or floors are like European-type options on T-bills or on eurodeposits or euroloans. You are, of course, thinking of calls as being relevant when prices are high, and puts when prices are low. But market values of loans are inversely related to interest rates. So a floor on the interest rate is a cap on the price and vice versa:

Example 16.3

Suppose that you have a one-year, NZD 100m loan, with half-yearly interest payments capped at 4.5 percent *p.a.*—that is, 2.25 percent per half-year. The interest rate for the period that starts immediately is already known—say, 4 percent *p.a.* The 2.25 percent cap on the next six-month effective return means that, after six months, you have the right to borrow NZD 100m at 2.25 percent (effective) for another six-month period. That is, six months from now you have the right to place (=sell) a new six-month promissory note with expiration value NZD 102.25m at a price of NZD 100m—a right that is valuable to you if at the reset date the six-month rate is above 4.5 percent *p.a.* and the normal market value of a six-month 102.25m note, therefore, below 100m. In standard optionspeak, you hold a put option on a NZD 102.25m note at a strike price set at $X =$ NZD $100m$.

DoItYourself problem 16.1

Suppose that you have a one-year, NZD 1m loan, with half-yearly interest payments with a floor at 3.5 percent *p.a.*—that is, 1.75 percent per half-year. The interest rate for the period that starts immediately is already known; for instance, it may be 4 percent *p.a.* or 2 percent effective. Interpret the 3.5% floor as an option on a PN: who holds and who writes the option, what type of option is it, what exactly are the terms and conditions of the underlying PN?

Both the example and the DoItYourself are summarized in Table 16.2.

In short, the floor on the interest rate is a call option on a promissory note, and the option is held by the lender and written by the borrower. The cap on the interest rate is a put option on a promissory note, and the option is held by the borrower and written by the lender. The reason why it is useful to re-state caps and floors as puts and calls on PN's is that a PN, unlike an interest rate, is an asset. So one can express a Put-Call parity in terms of option prices, underlying PN price, and discounted strike price. There is no similar direct link between option values and interest rates, except when prices are expressed as functions of interest rates:

DoItYourself problem 16.2

Buying a European-style call and selling a European-style put still means a forward purchase, and the forward purchase at strike X still has the same value as the underlying. So Put-Call parity still takes the form

$$\begin{pmatrix}\text{Call pre-}\\ \text{mium}\end{pmatrix} - \begin{pmatrix}\text{Put pre-}\\ \text{mium}\end{pmatrix} = \begin{pmatrix}\text{market value for-}\\ \text{ward purchase}\end{pmatrix} = \text{PV asset} - \frac{X}{1+r_{t,T}}. \quad (16.1)$$

Write the PV of the asset as a function of the limit rate (e.g. the 3% from the example) and the current market rate. Show that the right-hand side of the equation can be written as the *discounted* difference between the limit and market effective rates of return.

Hasty traders occasionally ignore the discounting, and express option prices as *p.a.* percentages so as to get a link with the p.a. interest rates in the formula.

Costs of a loan

There are various costs associated with a euroloan. These include:

- An up-front *management fee* and *participation fee*, sometimes 0.25 percent and sometimes a few percentages, see below. The up-front feature means that this amount is deducted from the principal. That is, the borrower receives only 99 percent to 99.5 percent of the nominal value of the loan.
- A *paying agent's fee* of a few basis points to cover the administrative expenses.
- The *risk spread* above the risk-free rate (that is, above LIBOR in the case of a floating-rate loan, or above the long-term fixed rate paid by a government of excellent credit standing). This spread depends on the quality of the borrower, the transfer risk of his or her country, the maturity and grace period, and the up-front fee. Also, the market situation affects the spread: there are strong

cycles, with spreads widening when something bad happens, then competition gradually narrowing the spread until a new bad event happens and so on.[9]

In principle, the fees are compensation for the services of the intermediaries, while the spread is a compensation for default risk. However, one can trade a higher up-front fee for a lower spread, and vice versa. For instance, borrowers often accept a high up-front fee in return for a lower spread because the spread is sometimes seen as a quality rating. One corollary of the trading of up-front fees for risk spreads is that the spread that country X pays may be a poor indicator of the creditworthiness of country X: an ostensibly reassuring spread may have been bought off by a large upfront fee. Another corollary is that reliable comparisons between offers from competing banks can be made only if there is a single, overall measure of cost. Thus, when comparing offers from competing syndicates, one should convert the up-front fees into equivalent spreads, or the spreads and paying-agent fees into equivalent up-front costs.

Example 16.4

Suppose that an up-front fee of USD 425,000 is asked on a five-year loan of USD 10,000,000 with an annual interest payment of 5 percent (including spreads) and one single amortization at the loan's maturity date. The effective proceeds of the loan are, therefore, USD 10,000,000 − 425,000 = USD 9,575,000. The effective interest rate can be estimated by computing the internal rate of return or yield, denoted by y, on the transaction:

$$\text{find } y\text{: } 9,575,000 = \frac{500,000}{1+y} + \frac{500,000}{(1+y)^2} + ... + \frac{10,500,000}{(1+y)^5}. \tag{16.2}$$

This equation can be solved on a spreadsheet or on a calculator. The solution is, approximately, $y = 6.0092$ percent, which is about 1 percent above the stated rate. Conversely, the up-front fee is equivalent to adding 1 percent *p.a.* to the stated rate.

In the above example, the future payments are known because the loan had a fixed interest rate. If the loan has a floating rate, one can no longer compute the yield because the future cash flows are unknown. However, the up-front fee can still be translated into an equivalent annual payment or equivalent annuity, using the interest rate on a fixed-rate loan with the same life and the same default risk. (To get the required number, simply ask a quote for a fixed-for-floating swap). The equivalent annuity can then be converted into an equivalent percentage spread by dividing the annuity by the loan's nominal value.

Example 16.5

We use the same data as in the previous example except that the loan has a floating

[9] The famous hedge fund Long-Term Capital Management (LTCM) was betting on a shrinking spread when, instead, a very bad thing happened, Russia's default. Betting again on a falling spread, LTCM was again wrong-footed, notably by the Asian crises. That was the beginning of the end for LTCM.

rate. If the normal all-in market rate on a fixed-rate loan with the same life and default risk as the floating-rate loan is 6 percent, the equivalent annuity (EqAn) of USD 425,000 up-front is determined as follows:

$$\begin{aligned}\text{EqAn: } 425,000 &= \frac{\text{eqAn}}{1+y} + \frac{\text{eqAn}}{(1+y)^2} + ... + \frac{\text{eqAn}}{(1+y)^5}, \\ &= \text{eqAn} \times 4.212367; \\ \Rightarrow \text{eqAn} &= \frac{425,000}{4.212367} = 100,893.47 \qquad (16.3)\end{aligned}$$

Thus, the up-front fee is equivalent to a spread of 100,893.47/10,000,000, that is, about 1 percent *p.a.*

If you ever have to do sums like this on a no-frills calculator, the shortcut to remember is

$$\frac{1}{1+y} + \frac{1}{(1+y)^2} + ... + \frac{1}{(1+y)^n} = \frac{1-(1+y)^{-n}}{y}. \qquad (16.4)$$

DoItYourself problem 16.3

If you would apply the approach of the last example to the one before, you would have found an equivalent spread of 1.0089 percent, not the 1.0092 of the earlier example. Why is there a difference? Which would you think the best figure (assuming anybody would bother about differences so tiny as this one)?

Credit Lines

In addition to outright loans, eurobanks also offer standby credits. These come in two forms:

- A standard line of credit (or credit line) of, say, GBP 100m gives the beneficiary the right to borrow up to GBP 100m, at the prevailing interest rate plus a preset spread. The difference between a credit line and a loan is that with a credit line the company is not forced to actually borrow the money: money is *drawn down* only if and when it is needed, and paid back at any date prior to the expiry date. Interest (in the strict sense) is payable only on the portion actually used, while on the unused funds only a *commitment fee* of 0.125 to 1 percent *p.a.* needs to be paid.

  A credit line is, in principle, a short-term commitment—say, three months. In practice, a credit line tends to get extended, but this is not an automatic right to the creditor. Unless stated otherwise, the credit line can be revoked by the bank if there are substantial changes in the creditor's credit standing.

- Under a *revolving commitment*, the creditor has the irrevocable right to borrow up to a stated limit, at the then-prevailing rate plus a preset spread during an agreed-upon period of (usually) several years. For instance, a borrower may have the right to borrow up to GBP 50m at interest of six-month LIBOR plus 1.5 percent *p.a.* This is similar to a credit line, except that it cannot be revoked during its life.

A credit line is like a single, short-lived option on the preset spread, and the revolving commitment is like a series of such options (one expiring every six months, for instance). These contracts are options, not forward contracts, because the beneficiary can always borrow elsewhere if the market-required spread drops. The credit line and the revolving commitment differ from caps in the sense that the contract imposes a ceiling on the spread, not on the interest rate.

Example 16.6

A company has the right to borrow at 1 percent above LIBOR. If the company's credit rating deteriorates, or if average spreads in the market increase, the 1 percent spread has become a bargain relative to what would have to be paid on new borrowing, and the credit will be effectively used. If, on the other hand, the company's rating improves, or if average spreads in the market fall, the 1 percent spread may be very high.

If the company uses the credit line, it still has to pay the agreed-upon 1 percent spread. However, the company can also borrow elsewhere, at a spread that reflects its better standing or the lower average spreads. Thus, the company has a cap option on the 1 percent spread.

This finishes our review of international banking products. We now describe their counterparts in the securities markets.

16.2 International Bond & Commercial-paper Markets

Almost simultaneously with the emergence of euromoney markets, firms and governments started issuing USD bonds outside the US, and sold the bonds to non-US residents. Such a bond was called a euro-dollar bond. As of the sixties, and particularly in the seventies, some eurobonds were denominated also in currencies other than the USD. Even though the dollar has long preserved its dominant market share, the fraction of dollar-denominated bonds occasionally drops below the total for European currencies, nowadays. Also in the seventies, corporations and governments started issuing short-term paper, although this short-term eurocommercial paper market did not really take off until the late eighties.

The markets to be discussed in this section are the tradable-security versions of the banking products that we discussed in the preceding section. Table 16.3 matches the eurobanking products with the closest equivalent in the eurosecurities markets. You may want to check these correspondences as we proceed.

Table 16.3: **Relationships between international banking products and securities**

Banking	Securities
Short-term	
Money Market	**Commercial Paper (CP) Market**
– short-term loan	– CP issue
– short-term credit line	– CP program
– rolled-over credit line	– note issuing facility (NIF)
– revolving commitment	– revolving underwritten facility (RUF)
– FF, FRA	– interest-rate futures
Medium- and Long-term	
Longer-term loans	**Notes and Bonds**
– fixed-rate loans	– fixed-rate bond
– revolving or floating-rate (FR) loan	– floating-rate note (FRN)
– FR loan with cap	– HIBO (higher-bound) bond
– FR loan with floor	– LOBO (lower-bound) bond

16.2.1 Why Eurobond Markets Exist

The explanations for the long-term success of international securities markets are largely similar to the ones cited for eurocurrency markets:

Lighter regulation for international public issues A bond issue aimed at the general public of one particular country is subject to many rules and regulations (although less so now than in the fifties and sixties). There are usually all kinds of publication requirements, and the issue has to be examined and approved by one or more regulatory agencies. In many countries, there are or were also issuing calendars (and, hence, queues) because the local government does not want foreigners to affect the country's reserves, money supply, or exchange rate; nor does the government want foreigners to "crowd out" local borrowers—especially not the government itself. By contrast, "international" transactions tend to be less regulated. For one thing, monetary authorities and capital market regulators are less concerned with issues that do not involve their own currencies and are targeting a few, well-off and well-informed investors or (even better:) foreign investors. This lack of regulation is especially true for tax-haven states that are often used as launching pads for international issues, but they also hold for *private issues* in mainstream countries.

Swift and efficient private placement By traditional US standards, publication requirements in Europe were never very stringent, and no rating is required for euro-issues. Even these comparatively lax requirements can be largely or entirely avoided if the issue is private rather than public. For loans privately placed with a limited number of professional investors, there is no queuing, and there are no (or almost no) disclosure requirements. In the EU, for instance, the telling feature is whether face values of EUR 50,000 or less are being offered; if so, the issue is deemed to be targeted at retail investors rather than professionals, and a prospectus must be published, approved by the local Central Bank. (Approval by one Central Bank suffices to sell anywhere in the Union—the so-called "passport".) But large-denomination bonds escape all this hassle.

Simple contracts As borrowers are generally of good credit standing, eurobonds tend to be unsecured; thus, legal costs, as well as the expenses of bonding and monitoring, are avoided. Since lenders have no control whatsoever over the borrower, only well-reputed companies can play this game at a reasonable price; small players often find the risk spread they would have to offer unattractive.

Tax games Eurobonds, being anonymous bearer bonds, traditionally make it easier to evade income taxes. Withholding taxes can be avoided by issuing the bonds in tax havens, and most OECD countries have recently waived withholding taxes for nonresidents.

Large issues Issues below USD 100m are very rare, nowadays. Most issues are now 500-1000m USD or EUR, but even 5b issues are placed in a day or two (not including the unofficial bookbuilding) and no longer raise eyebrows. In 2006, the largest issue was 22b. Such a big placement allows relatively low issuing costs.

Disintermediation Since the mid-seventies, impetus for the growth of the eurosecurities market has come from the general *disintermediation* movement. Disintermediation means "cutting out the intermediary"; that is, corporations borrow directly from investors. This evolution was the result of two forces. First, in the 80s, many banks lost their first-rate creditworthiness when parts of their loan portfolios turned sour (due to the international debt crisis[10] and the collapse of the real estate markets[11]) As a result, these banks were no longer able to fund at the AAA rate, which meant that top borrowers could borrow at a lower cost than banks could—by tapping the market directly. Second, as a response to the lower profits from lending and borrowing and to the stricter capital adequacy rules, banks preferred to earn fee income from bond placements or commercial paper issues. Unlike operations involving deposits and loans, this commission business creates immediate income (rather than income from bid-offer spreads, received later on) and does not inflate

[10]Emerging-market debt had ballooned after, in 1974, the oil price doubled (which made many countries borrow heavily) and when a wave of inflation in the late 1970s had increased interest rates to unusual levels (much of the oil debt was at floating rates). Borrowers defaulted or renegotiated both their bank debt (in the "London Club") and their Government-to-Government debt (in the "Paris Club").

[11]One background item was that, in the early 1970s, the distinction between thrifts and banks was lifted. Thrifts (or Saving&Loans) were originally cooperatives where members made time deposits and got time loans, mostly mortgage loans. Unlike commercial banks, they could not take sight deposits and give overdraft facilities. When the distinction was lifted, the old S&L were left with far less controls than commercial banks. As a result they made many bad investments, contributing to a boom and bust in real estate. The mess took years, and trillions, to sort out. The real-estate bubble also spread to Europe and hurt also the old commercial banks, most notably in the US, the Nordic countries and France.

the balance sheet.

16.2.2 Institutional Aspects of the International Bond Market

We briefly describe some institutional aspects of the international bond market.

Bearer securities Eurobonds are bearer bonds, that is, anonymously held rather than held by investors listed in a register. In the old days, "bearer" actually meant "made out to bearer"—actual pieces of paper, with coupons that can be clipped off and cashed in by the holder. The principal of the bond was represented by the mantle, the main part of the paper (after the coupons have been clipped off). In many countries, an investor can cash in coupons and principal paid out by bearer securities without having to reveal his or her identity to the bank that acts as paying agent. In contrast, if the security had been a registered security, the issuer would know the identity of the current holder of each bond, and pay interest by mailing a check. US domestic bonds are usually registered, nowadays. In the EU, even bearer bonds tend to be *non-deliverable* nowadays, that is, not physical pieces of paper; investors buy them electronically from intermediaries, but the issuer still does not keep a register.

Interest payments Eurobonds originally carried (and to a large extent still carry) fixed coupons. Coupons are most often paid annually instead of every six months (the domestic US pattern). Floating-rate notes (FRN) gained popularity when interest rates rose, in the 1970s-80s, making many borrowers hesitant about long-term fixed-rate bonds; when interest rates are low, in contrast, investors tend to be the party that shuns fixed interest rates. In a floating-rate loan, the interest rate is periodically reset on the basis of the then-prevailing LIBOR for that horizon plus a preset spread. Sometimes, the FRN has a cap or floor on the floating interest rate. Capped FRNs are sometimes called *HIBO bonds* (*higher-bound bonds*), and floored FRNs *LOBO bonds*. *Perpetual FRNs* were briefly fashionable in the mid-eighties.

Amortization Amortization of the bond's principal amount typically occurs at maturity. Such bonds are known as *bullet bonds*. Alternatively, the borrower may undertake to buy back predetermined amounts of bonds in the open market every year. This is called a *purchase-fund provision* or a *sinking-fund provision*. Under a variant provision, the borrower does not have to buy back the bonds if market prices are above par. Instead, the borrower has a right to call a predetermined part of the issue every year.

Currency of denomination The currency of denomination of the bonds is most often a single currency (especially the USD, DEM or, now, EUR, JPY, and CHF). Also the private ECU gained some popularity as a currency of denomination in the early-to-mid 1990s. Other currency baskets, such as the SDR or the European Unit of Account, have never really caught on. Some bonds have currency options attached to them. Such currency options bonds are discussed in Chapter 8. Occasionally, you also see a dual currency bond, which pays out its coupons in one currency and the principal in another currency.

DoItYourself problem 16.4

Suppose that the holder of a five-year bond receives an annual coupon of USD 500 and can choose to receive at maturity either USD 10,000 or EUR 10,000. Taking the USD as home currency, you can describe this bond in two ways, one involving a put and one involving a call. Find these two descriptions, and link them via Put-Call Parity.

Stripped bonds Bond stripping essentially means that the coupons and the principal components of the bond are sold separately. If bonds are actual pieces of paper made out to bearer, you can strip bonds at home with a pair of scissors: just clip off all of the remaining coupons, and sell them separately from the mantle, the piece that stands for the principal. On a larger scale, and especially when bonds are registered rather than bearer securities, stripping is done by buying coupon bonds, placing them into an incorporated mutual fund or a trust, and issuing separate claims against this portfolio, representing either the coupons or the principal.

The main consequence of stripping is that the principal can be sold separately, as a zero-coupon bond. One motivation for stripping is that immunisation and asset/liability management are simplified if there are zero-coupon instruments for many maturities. Also, zero-coupon bonds, offering capital gains rather than interest, get favorable tax treatment in many countries. In some countries, including Japan and Italy, capital gains are often partially exempt from personal taxation. Thus, the principal is sold to *e.g.* Japanese or Italian investors, and the (taxable) coupons are sold to low-tax investors.

Issuing procedures (1): the consortium Placement of eurobonds is most often through a syndicate of banks or security houses.

- The *book runner* (formerly called *lead bank* or *lead manager*) negotiates with the borrower, brings the syndicate together, makes a market (at least initially), and supports the price during and immediately after the selling period. Book runnership can be shared by a group.
- There are often, but not always, *managing banks* that underwrite the issue and often buy part of the bonds for their own account.
- The *placing agents* call their clients (institutional investors or individuals) and sell the bonds on a commission basis.

  Just like in the case of bonds, there is creeping title inflation. More and more often the underwriters are called lead managers, and the term co-managers then refers to firms that just distribute the paper (the placing agents of old).
- The *fiscal agent* takes care of withholding taxes, while the *trustee bank* monitors the bond contract (if any such contract exists; most bonds are unsecured and do not have bonding clauses).

The various players get their commission through the discount they get when they buy the paper. In the table below we start from a set of commission percentages and then work out their implications for the prices at which the players buy and resell the bonds. The paper is assumed to have a nominal value of 10,000 per unit.

	% commission specs	the bank buys at	... and sells at	
lead manager	0.5%	9,750	9,800	(to underwriters)
underwriters	1.0%	9,800	9,900	(to selling agents)
sellers	1.0%	9,900	10,000	(to public or back to underwiters)

Prospective customers can find information about the issuing company and about the terms and conditions of the bond in a prospectus. Often, an unofficial version of the prospectus is already circulating before the actual prospectus is officially approved. This preliminary prospectus is called the *red herring* and is part of the *bookbuilding* stage, where the putative managing group is gauging the market's willingness to buy. Once the decision to issue has become final and the prospectus is official, investors can already buy forward the bonds in the few weeks before the actual issuing period starts. This period of unofficial trading is called the *gray-market period.*

The whole process typically takes up about a month or more—not exactly fast. For this reason, alternative issuing procedures have emerged:

Issuing procedures (2): alternatives to the consortium One rarer solution is the *bought deal*: a bank single-handedly buys the entire issue, before building a book and finding co-underwiters. This is riskier to the bank, so typically the implied underwriting fee is larger and/or the issue smaller—one always pays some price for speed.

Other alternatives entirely omit underwiting. In a *fixed-price reoffer*, the price to be paid by the public is set, and sellers get a commission if and when they place paper, say 0.15 percent. The borrower bears the risk that the whole issue flops, but since the procedure is faster than the traditional consortium method, the risk is thought to be bearable, by some. Still, one rarely sees such deals. In a *yield pricing* issue the price is set at the very end, taking into account yields of comparable bonds in the secondary market. The issuer agains buys speed, and there is far less risk that the paper is unsellable, relative to the case where the coupon is set weeks before the actual placement. But rates can still change after the selling starts, or the risk spread may not please the market, so there obviously is no certainty about quantity and price until the selling is over.

Another method is like the traditional *au robinet* (on tap) method, the way a bank traditionally issues its own retail CDs to the general public. This is best known as the *Medium-Term Note* (MTN) method even though it is now used for paper of 9 months to up to 30 years. Here, the borrower mandates a bank to sell paper within certain guidelines, say "money in the 1-5 year range at 50 basis points or

Figure 16.2: **EBRD Ruble bond Issue**

Nine banks underwrite new EBRD rouble bond

Rate on first quarterly coupon set at 5.56 percent

The EBRD has completed the placement of its second rouble bond, underwritten by a syndicate of nine international and Russian banks. The rate on the first coupon has been set at 5.56 percent.

This new 5-billion rouble (equivalent to €147 million) five-year floating rate instrument is being launched by the EBRD to meet the strong demand in Russia for the Bank's local currency loans. The EBRD raised its first rouble bond in May 2005 for exactly the same amount and with the same tenor.

The Russian subsidiaries of Citibank and Raiffeisenbank Austria are the Joint Lead Arrangers of the new issue – with JP Morgan Bank International, ABN Amro Bank AO, ING Bank (Eurasia) ZAO, Bank WestLB Vostok (ZAO) acting as senior co-lead managers. SAO Commerzbank(Eurasija) and Gazprombank are co-lead managers. Vneshtorgbank is the co-manager. ING will also act as the Calculation Agent for the issue.

The new EBRD bond's floating rate coupon is once again linked to MosPrime Rate, a money market index launched last year under the auspices of Russia's National Currency Association (NCA).

The MosPrime rate is calculated daily for 1-months, 2-months and 3-months deposits based on the quotes contributed by eight banks: ABN Amro, ZAO Citibank, Gazprombank, International Moscow Bank, Raiffeisenbank, Sberbank, Vneshtorgbank and WestLB.

The launch of the EBRD's second floating rate note underscores the development of the MosPrime Rate as a widely accepted money market benchmark in Russia since its launch in April 2005. Several public transactions have been linked to this index in the past year, the most recent, as well as the largest, being a 7.2 billion rouble (equivalent to €212 million) loan for Mosenergo.

The new issue was registered with the Federal Financial Markets Service (FFMS) on April 11. Just as with the first issue, the EBRD will apply for its bonds to be listed and traded on the Moscow Interbank Currency Exchange (MICEX) and for the Central Bank to include them in its Lombard list. This would make the bond available for repo transactions with the Central Bank.

The issue pays a quarterly coupon, with the coupon rate reset every three months in line with the then prevailing MosPrime offered rates. The coupon rate for the bond will be published at Reuters page EBRDRUBFRNRATE.

The EBRD enjoys an AAA/Aaa rating from international rating agencies.

Press contact: _Richard Wallis, Moscow - Tel: +7095 787 1111; E-mail: wallisr@ebrd.com

Terms and conditions Sitemap Feedback

Source http://www.ebrd.com/new/pressrel/2006/40apr28.htm

less over the relevant US T-bill rate"; and the mandatee simply waits for queries from big investors with excess liquidities. A deal like this can be made quite fast, occasionally even within one hour, and costs are quite low. One reason is that there is no official solliciting, no prospectus etc is needed. In addition, the intermediary is not guaranteeing anything. But obviously the issuer has no idea how long it will take to raise the sum they had in mind, and may have to improve the terms after a while. Still, this issuing procedure has become a serious alternative to the consortium system.

Secondary market The secondary market for eurobonds is not always very active. Many bonds are listed on the Luxembourg Bourse, but this is largely a matter of formality. A few hundred issues trade more or less actively on London's International Stock Exchange Automated Quotation (SEAQ International) computer system. Through SEAQ International, market makers post bid-and-ask prices for non-UK blue-chip stocks and for eurobonds. There is also an over-the-counter market, where (bored) bond dealers keep buying and selling to each other. Multilateral clearing institutions like Euroclear in Brussels, Clearstream (formerly Cedel) in Luxembourg, and the London Clearing House reduce the costs of physical delivery of the bond certificates themselves. (They also offer clearing services for trades of stocks listed on exchanges; Clearstream is now owned by Deutsche Börse.)

Figure 16.2 shows a press release by the European Bank of Reconstruction and Development on a Ruble bond issue.

Figure 16.3: **A RUF: Kertih Terminals**

Kertih Terminals Signs RM500 Million Financing Agreement For Bulk Chemical Storage Project

Kertih Terminals Sdn Bhd (KTSB) has signed an agreement with RHB Sakura Merchant Bankers Bhd as Arranger and Agent for a RM500 million Revolving Underwritten Facility (RUF) with Term Loan Conversion to finance the development of its centralised liquid bulk chemical storage and handling facility in Kertih, Terengganu. The financing agreement was signed today in Kuala Lumpur between KTSB, RHB Sakura Merchant Bankers and a group of financial institutions as underwriters and tender panel members of the RUF.

Under the agreement, KTSB, taking advantage of the prevailing favourable interest rates, will issue short-term negotiable debt instruments directly to investors during the first five years of the RUF, after which the facility is convertible into a four-year Term Loan. Additional features of the RUF, which has been assigned a short-term rating of MARC-1 by Malaysian Rating Corporation Bhd, include the option to raise fixed rate debts via conventional borrowings, structured debts (bonds) instruments or Islamic financing instruments.

KTSB, a joint venture between PETRONAS (40 percent), GATX Terminals (Pte) Ltd (30 percent) and Dialog Equity Sdn Bhd (30 percent), is undertaking the centralised chemical storage project which forms an integral part of the Kertih Integrated Petrochemical Complex (IPC) currently being developed by PETRONAS. Phase one of the storage project is at an advanced stage of construction. When fully operational, the facility will have 37 tanks with a total storage capacity of 403,358 cubic metres to cater to a host of users and customers at the Kertih IPC. These include Vinyl Chloride (Malaysia) Sdn Bhd, PETRONAS Ammonia Sdn Bhd, BP PETRONAS Acetyls Sdn Bhd, Aromatics Malaysia Sdn Bhd and the Union Carbide Corporation-PETRONAS' derivatives joint venture.

Issued by:
Kertih Terminals Sdn Bhd
109 Block G, Phileo Damansara 1
No 9 Jalan 16/11
46350 PETALING JAYA
Tel: 03-7551199

Source http://www.petronas.com.my/internet/corp/news.nsf/2b372bb45ff1ab3a48256b42002b19a7/a09ff4bccca787b048256adf0049a50f?OpenDocument

Eurobonds represent the long end of the eurosecurities market. We now turn to markets for securities with shorter times to maturity.

16.2.3 Commercial Paper

Commercial paper refers to short-term securities (from seven days to a few years) issued by private companies. Just as bonds are the disintermediated version of long-term bank loans, commercial paper (CP) forms the disintermediated counterpart of short-term bank loans. CP markets have existed in an embryonic form ever since banks drew promissory notes on their borrowers as a way to document loan agreements. However, the market became important only in the eighties when, as part of the general disintermediation movement, large corporations with excellent credit standing started issuing short- and medium-term paper, which then was (and is) placed directly with institutional investors. The volume of the market remains low relative to the bond and bank-loan market.

The market consists of notes, promissory notes, and certificates of deposits (CD)s. *Notes* are medium-term paper with maturities from one to seven years, usually paying out coupons; many Europeans would simply call them bonds. *Promissory notes* have shorter lives (sometimes as short as seven days), and are issued on a discount basis, that is, without interim interest payments. Notes and promissory notes issued by banks are called *certificates of deposit* (CDs).

Although an CP issue can be a one-time affair, many issuers have an CP-program contract with a syndicate. A bare-bones CP program simply eliminates the bother of getting a syndicate together each time commercial paper needs to be placed, but many programs also offer some form of underwriting commitment (for issues up to a given amount and within a given period). Such a commitment can stipulate the following terms:

A fixed spread An arrangement under which a borrower can issue CP at a fixed spread, e.g. 0.5 percent over LIBOR, is called a *Note Issuing Facility (NIF)*. This preset spread may become too high later on, notably if the borrower's rating improves or if the average spread in the market falls. In such cases, the borrower loses—he pays too much, in view of the changed circumstances—and the placing agent gains because he or she can place the paper above the initially anticipated price. In contrast, if the preset spread becomes too low, the borrower unambiguously wins; the cost is then born by the underwriter, who has to buy the issue at a price that exceeds the fair market value.

Figure 16.3 refers to a RUF extended to Malaysia's Kertih.

The difference between a NIF and a RUF is less important than what it may seem at first. Even a NIF is an option on a spread, not a forward contract on a spread, because the borrower is under no obligation to use the facility. That is, if spreads go down in the market, the borrower can simply forget about the NIF and issue paper through a new syndicate or under a new agreement. Under such circumstances, the advantage of the RUF to the borrower is that it avoids the cost and complications of setting up a new issuing program and, of course, that there is a cap on the risk spread.

16.3 How to Weigh your Borrowing Alternatives

Companies can get tentative offers from more than one bank or group of banks, or offers in many currencies. How to compare them?

One of this book's fundamental tenets is that, in a perfect market, everything is priced correctly, and nothing is gained nor lost by the mere switching from one borrowing alternative to another. NPV's on all financial transactions would be zero, in the sense that the PV of the future service flows is fully reflected in the price. Yet this does not mean that a real-world CEO can always relax and pick a loan at random from any first-coming bank. Let's review a few arguments that came up already in the preceding chapters, and add a few new ones. We group the relevant items under two headings: interactions with operational cash flows, and market imperfections.

Interactions with operations To a company, a financial contract may deliver more than the contract's very own cash flows; notably, as we saw in Chapter

12, the choice of the denomination of one's assets or liabilities may interact with the operational cash flows in wider sense, for instance, by affecting the probabilities of financial distress and the costs that come with it. If so, these interactions would affect the firm's market value.

Imperfections Many aspects of real-world markets could make the choice between borrowing alternatives relevant. Information asymmetries among lenders are likely to lead to inconsistent risk spreads across banks, for instance, or tax asymmetries may make high-interest or high-volatility currencies more attractive. An even more fundamental imperfection would be if prices in exchange and money market are fixed by the government and/or if a license *raj* prevents arbitrageurs and speculators to do their usual job: then even deals at the risk-free rate, assuming away any information or tax issues, are likely to come with non-zero NPV's. Non-zero NPV's could also arise from less glaring forms of market inefficiency, though, like herd behavior—anything that might lead to mispricing, which the astute speculator can then take advantage of. Lastly, there are the fees that banks charge, and the careful money manager has to check and re-check that the lenders are not trying to overcharge.

We start with the issues that should be the most likely cause of relevance in well-developed markets: costs and risk spreads. These are also easily quantified and summarized into one number. Having ranked the alternatives in terms of these items, we can then assess whether there is a good reason to deviate from that ranking. The easiest case is one where the home and foreign capital markets are both very open and developed. Agents can freely chose where to borrow, from whom and in what currency they want. There are, in addition, competitive swap markets where foreign-currency borrowing can be separated from borrowing abroad. In such a situation it is plausible that, if there were no default risk and no information asymetries, little value would be gained or lost by switching to another currency—whether we do so explicitly or via a swap. In such a setting we can focus on just the costs asked by competing banks over and above the risk-free rate, that is, the items reflecting default risk and information asymmetries.

16.3.1 Comparing all-in Costs of Alternatives in Open, Developed Markets

Let's work via an example. The issuer is a US company that has a preference for USD borrowing; but there is a EUR offer too. The hoped-for proceeds would be USD 200m before costs, or, at the spot rate of USD/EUR 1.25, EUR 160m, and the CFO is going for a 7-year bullet loan. Table 16.4 shows the conditions, along with some other useful data and the computations. Please refer to them as the discussion proceeds.

Evaluation under Idealized Circumstances

We could look at the sum of discounted risk spreads and other costs, using the swap rate as the discount rate. This is similar to what we did in Chapters 5 (on forward contracts) and 7 (on swaps), except, trivially, that now we add an upfront cost. But there we took the swap dealer's point of view, whose risk is not the same as those born by a lender. So let's first bring a story that tells us why a procedure like this also makes sense, subject to a caveat, for lenders and borrowers without right of offset.

We regard a bond issue or a bank loan as an NPV problem. To the borrower, the proceeds are immediate, and the costs are the subsequent service flows—just the reverse of what one sees in capital budgeting, but that is not important. Let us denote the swap rate, a risk-free yield-at-par, by s; the spread as asked by banker b over and above s by ρ_b; and the required discount rate by R. Finally, let V_{nom} denote the gross size of the loan, and U_b the upfront cost proposed by banker b.

The NPV of accepting b's proposal equals the net proceeds, $V_{nom} - U_b$, over and above the PV of the future service streams. We write this in line one, below. In line two, we just simplify and regroup, as follows:

$$\begin{aligned} NPV_b &= \underbrace{V_{nom} - U_b}_{\text{net IN}} - \underbrace{V_{nom} \times (1 + (s + \rho_b - R) \times a(R, \#years)}_{\text{PV of subsequent outflows}}, \\ &= V_{nom} \times (s - R) \times a(R, \#years) - \underbrace{[U_b + V_{nom} \times \rho_b \times a(R, \#years)]}_{\text{the bank-specific part}} \end{aligned} \tag{16.5}$$

Offers may have been asked from various bankers, all for the same amount V_{nom}; and s and R are market-wide numbers. Thus, the first part in the NPV expression, Equation [16.5], is common to all offers, and for that reason the competing offers from various banks can be ranked by looking just at the second part, the upfront cost plus the PV'ed spreads, labeled "bank-specific part" in the equation. The spread asked consists of the "objective" risk premium one would see in perfect markets, plus, in realistic bond markets, a compensation for the investors' unfamiliarity with the borrower: unknown parties look more risky. The investment bank, in the case of a bond issue, tries to reduce the unfamiliarity premium by road shows etc, but this increases U_b, the upfront cost. In addition, part of the bank's upfront expenses may be paid for not via the upfront fee U_b but by an extra interest spread instead; alternatively, as we have seen, the parties may agree to lower the very visible spread, and increase the less visible upfront fee U_b instead. That's why we should look at the whole package. Thus, we propose the PV criterion,

$$\text{PV'ed total bank-specific component } = U_b + V_{nom} \times \rho_b \times a(R, \#years). \tag{16.6}$$

This is what we did when we compared risk spreads in the Swaps chapter, except that we add the upfront cost and we use R not s. Using s was justified for a swap dealer, who benefits from the right of offset and the credit trigger. For a bank (or the counterpart, the lender), the uncertainties are usually greater, so very often one needs to be a bit more careful about default risk.

How well do we know R, in reality?

This looks like a cut-and-dried problem. The only hitch is that R, the required rate of return, isn't easily observable. Only the Great Banker In The Sky knows it well. How come? Can't we just take the prior that the NPV is zero, which would allow us to infer R as the internal rate of return? In perfect markets, of course, NPV's from financial transactions must be zero. In reality we cannot bank on that, though, because acquisition of information is costly and time-consuming. This is especially an issue in the case of risky corporate borrowing, which is full of information asymmetries—either between banks and borrowers, or among the banks that might compete to act as lenders. Let's look at each of these asymmetries.

- If the financiers know more about the market situation than you do, chances are that they make a gain and you a loss. Not *all* bankers are angels. There is court evidence how investment bankers have underpriced the IPOs they managed, so as to be able to dole out goodies to friends and cronies. You may also have heard how derivatives dealers openly mailed each other about the "rip-off factors" they had included in their contracts, and how during the IT bubble investment bankers made fun of the "fools" they sold to.

- These are, of course, just anecdotes; but there exists a respectable academic literature on "hold-up" behavior. House bankers have a bit of a monopoly position, so the argument goes, since they have built up long-term knowledge about the borrower. Breaking up the relation would be costly for the borrower, since it takes time and effort for another bank to just re-discover all the info and insights the incumbent already has. Thus, the house bank is in a position to exact a monopoly rent—not too much, of course, otherwise they lose the account. Empirical evidence shows that banks actually do so.

For these reasons, mildly negative NPVs are far from unlikely. The borrower might still go along with a negative-NPV bond deal if, as pointed out, the loss is not large enough to justify changing banks or consortia and if the loss is small relative to the NPV of the direct investment that is being financed. In a way, the bankers just grab a slice of the firm's business gains. But the bottom line surely is that you cannot just postulate that competition is perfect, NPV's therefore zero, and R visible as the IRR of the deal.

Evaluation under Realistic Circumstances

Bearing all this in mind, let's now critically review two feasible methods and see how they relate to the ideal solution we just outlined.

Table 16.4: **Appalachian Barracuda Corp's Analysis of Funding Offers**

1. The competing Offers

	(a) swap rate	(b) loan rate	(b)–(a) spread	upfront fee	7-year annuity factors at swap rate	at IRR
USD	4.50%	4.00%	0.50%	2.00%	6.002 054 67	5.819 282 33
EUR	4.35%	3.80%	0.55%	1.75%	6.046 667 84	5.860 778 76

spot rate USD/EUR 1.25

2. Comparing the loans via swap-rate-based PVs

USD loan:

risk spreads (PV)	$200m \times 0.005 \times \overbrace{6.00205467}^{\text{a(7yrs, swap\$)}} =$	USD	6,002,055
upfront	$200m \times 0.02 =$		4,000,000
total cost		USD	10.002,055

EUR loan:

risk spreads (PV)	$160m \times 0.0055 \times \overbrace{6.046667843}^{\text{a(7yrs, swap€)}} =$	EUR	5.321,068
upfront	$160m \times 0.0175 =$		2,800,000
total cost		EUR	8.121,068
id. in USD:	$8.121{,}068 \times 1.25 =$	USD	10.151,335

Extra cost of EUR = USD 10.151,335 –10.002,055 = USD 0.149m

3. Comparing the loans via IRR-based PVs

USD loan:

$\text{YIELD}(\overbrace{\text{"1/1/2001"}}^{\text{some start date}}, \overbrace{\text{"12/31/2007"}}^{\text{7 years later}}, \overbrace{0.045}^{\text{coupon}}, \overbrace{98}^{V_t}, \overbrace{100}^{V_T}, \overbrace{1,1}^{\text{don't ask}}) =$		% 4.844
swap rate		% 4.000
all-in spread		% 0.844
cost in USD: $200\text{m} \times [0.00844 \times \underbrace{5.819\,282\,33}_{\text{a(7yrs, IRR\$)}}] =$	USD	9.823m

EUR loan:

YIELD("1/1/2001", "12/31/2007", 0.0435, 98.25, 100, 1, 1) =		% 4.649
swap rate		% 3.800
all-in spread		% 0.849
cost in EUR: $160m \times [0.00849 \times \overbrace{5.860\,778\,76}^{\text{a(7yrs, IRR€)}} =$	EUR	7.961m
cost in USD: $1.25 \times 7.961 =$	USD	9.951m

Extra cost of EUR loan = 9.951m–9.823m = USD 0.128m

Key Method 1 computes the PV-ed spreads using the swap rate s and then adds the upfront. The resulting cost difference is USD 149K. Method 2 computes an internal rate of return (IRR)—I show the spreadsheet command that does it for you—and finds the IRR is 0.844% above the swap rate for the dollar loan, and 0.849% for the euro loan. The cost above the swap rate is then PV'ed at the IRRs.

1. **PV the spreads at the swap rate; add upfronts** This is close to our first criterion, Equation [16.5], except that we use the swap rate s instead of the risk-adjusted rate R. In defense of this method, remember that we do not want to value a given loan; rather, we want to rank two loans on the basis of the difference of the cost components, over and above the swap rates. Thus, first, we only discount the bank-specific part; so most of the service streams are not considered, which eliminates also most of the valuation errors created by using s instead of R. In fact, the PV of the spreads ρ_b is mostly affected by the size of ρ_b, in the numerator, not so much by the discount rate. Second, we make the same mistake for all loan alternatives, so that the net impact on the calculated cost differential is even smaller.

2. **Compute an IRR, subtract the swap rate, and pv the total cost at the IRR** The IRR, familiarly, is the stand-in discount rate that would equalize the discounted value of the future payments to the net value (after upfront costs). (This must be done numerically; the table shows a spreadsheet command that provides the answer.) So this method simply postulates that the deal's NPV in Equation [16.5] is zero, which, if true, allows you to compute an estimate of R. This allows us to estimate a total-cost spread that can be discounted at the IRR.

   Assuming a zero NPV is not a crazy idea: in the absence of asymmetries it would actually be quite natural that both lender and borrower made a break-even deal. So this estimate of R must be close to the mark for big lenders with little information asymmetries. For smaller borrowers, negative NPV's are far more likely, in which case R is overestimated and the PV'ed cost underestimated.

   Example 16.7

   Think of the one-period case where we easily see what's going on. The swap rate is 8%. Suppose the fair value of a 10% loan is 100 but you are ripped off and get 99 only. The IRR would be 110/99–1 = 11.11% while the true R is 110/100–1 = 10%. Using the IRR we'd estimate the cost at (11.11–8)/1.11 = 2.799 while the true figure is 2/1.1+1=2.818.

The reassuring finding, in Table 16.4, is that the two measures of the differential cost are very similar. Using swap rates we'd reckon the cost difference between the USD and the EUR offers is USD 149K in favor of the HC offer, while the estimate is USD 128K when we use IRRs. The disagreement is 21K, a tiny number relative to the face value, USD 200,000K. Even more important, both methods agree that the HC loan, USD, has the lower costs.

A translated or equivalent spread for FC loans

In the above, I recommend that you size up the whole package in PV terms, an amount of cash money. Bankers and CFOs often look at percentages, though. Why

Table 16.5: **Percentage total spreads of borrowing alternatives**

	USD loan	EUR loan	difference
Using swap rates			
(a) PV'ed cost in USD	10.002	10.151	
(b) Eq. annuity: PV/6.002 054 67	1.666	1.691	
(c) id/200,000,000	0.833%	0.846%	0.012%
Using IRR*s*			
(a) PV'ed cost in USD	9.823	9.951	
(b) Eq. annuity: PV/5.819 229 22	1.688	1.710	
(c) id/200,000,000	0.844%	0.855%	0.011%
difference of % cost estimates	0.011%	0.009%	

PVsPVs instead? First, basic financial logic tells us to generally trust PVs, *i.e.* numbers in dollars or euros or pounds, not percentages: we pay for our shopping with money not percentages, and 1% extra on 1,000,000 means more money than 2% extra on 100. True, in this particular instance this is no issue since the alternatives, by design, all have the same scale, USD 200m, implying that there is no harm in looking at the percentages here. But there is a second reason why PVs, in units of money, are better than spreads: amounts of money are easily understood and easily compared across currencies. In contrast, to many business people it would be hard to see whether a spread of 0.75% on top of a swap rate of 8% is better or worse than a spread of 0.30% on top of a swap rate of 2%.

Despite all this, some traditional finance *babu*s insist on percentage spreads. If your boss really presses you, here is how you could respond without giving up rigor.

The simple way to come to a *p.a.*-spread type number is to divide the HC PV numbers by the HC annuity factor, which means that we compute the equivalent annuity of all costs, upfront or not. Then we express the equivalent annuity as a percentage. Table 16.5 shows the results. Note how, by always using the same number —HC annuity factor times HC face value—to rescale the PV'ed costs we cannot possibly change the ranking of the alternatives. It is, in fact, easily shown that the FC costs are those of swapped FC loans, not of the original FC loans. Just hope that your boss does not raise the question; and if (s)he does, say that the numbers are adjusted for currency risk.

The calculations show, first, that the estimated total spreads are not very much affected by whether you use swap rates or IRRs: these intra-column differences, shown in the bottom line, amount to one basis point only. The second conclusion is that both methods agree that the EUR and USD offers are very close, with a disadvantage of slightly over one basis point for the EUR loan. These differentials across columns are shown in the rightmost column of the table.

Making a decision

If cost is the only consideration, then in this example the USD offer has the edge, but it is a very close race. What other considerations could have swayed the balance? Basically, anything that would imply a preference for EUR would do, given that costs are essentially the same.

One consideration that could interfere with this conclusion would be speculation, the way we defined it in earlier chapters. The calculations here would be very different: instead of comparing the USD loan to the swapped EUR loan we'd have to consider the unswapped version and see whether the difference of the IRRs is justified by the expected currency movements. In Table 16.4 the IRR of the unswapped EUR loan was found to be 4.65%, against 4.84% for the USD offer. So if the EUR appreciates by less than 0.2% per year, on average, then in terms of expectations it would be less expensive than dollar borrowing. Early 2008, many may feel that the Euro is actually overvalued and is expected to slide back to lower levels. If, in addition, the yield is lower, then we'd have an argument for EUR borrowing. This logic is very different from the cost-based calculations, where any discrepancy between the differential swap rates and the expected rate of appreciation is postulated to be rational—for instance, reflecting risk considerations.

Speculation is not the only argument that might affect the final decision. There may be EUR-related assets that need to be hedged anyway. Bear in mind, though, that the existence of foreign assets does not necessarily mean that these assets come with a positive exposure; remember the Android example in Chapter 13.

You will agree that weighing the speculative and hedging aspects is difficult, as neither is easily quantified. But things are even less satisfactory when the foreign currency under consideration lacks financial instruments like forwards and swaps or, even worse, the foreign money and exchange markets are plagued by controls.

16.3.2 Comparing all-in Costs of Alternatives in Regulated, Incomplete Markets

The alternative to the USD loan (200m, as before) now is a CNY one, as the investment now is in China. There are no long-term forwards or swaps. There is no liquid government-bond market, and if there were one there still is the problem that there are exchange controls. Neither Chinese investors nor foreigners can freely switch their funds from CNY to USD, so that one cannot just assume that Yuan and Dollar loans are correctly priced relative to each other in one international market.

The hoped-for proceeds of a possible Yuan loan, at the spot rate of CNY/USD 8.00 (this is a rounded number to simplify the figures), would be CNY 1.600m. The CFO is still going for a 7-year bullet loan. The terms offered are a loan at 6.75% and total upfront fee of 1 percent. To see whether this is good or bad, you could look at the bids and asks of the People's Bank, the market leader, as shown in

Table 16.6: **People's Bank of China interest rates, late 2006**

lending rate		savings rate	
		3 months	1.80%
6 months	5.58%	6 months	2.25%
6-12months	6.12%	1 year	2.52%
1-3 years	6.30%	2 years	3.06%
3-5 years	6.48%	3 years	3.69%
> 5 years	6.84%	5 years	4.14%

Table 16.6. The lending rates are, in fact, the People's Bank's reference rates. The central bank is, indeed, the entity that decides on reference interest rates for loans, but it also sets a band within which local bank branches have some discretion in adjusting their lending rates. As of 1999, the band is 10% below and 30% above the reference rate and for loans to large enterprises, the upper limit is 10% above the reference rate. According to Garca-Herrero et al. (2005), due to the banks' lack of expertise in accessing borrowers' credit risk, most loans are just contracted at, or even below, the PBC's reference rate despite the additional flexibility provided by the liberalization of interest rates.

An uncritical spreadsheet-adept could calculate that the borrowing rate for best-quality borrowers would be the rate at the bottom of the People's Bank admissible range, *i.e.* 6.84% × 0.90 = 6.156%. The corresponding risk spread would then be 6.75% - 6.156% = 0.594%. Below, then, are shown the calculations using that stand-in for the risk-free rate; the corresponding annuity factor is 5.551,560,664. I show the NPV and total-spread calculations:

Example 16.8 CNY alternative

risk spreads (PV)	$(1600 \times 0.00594 \times 5.551,560,664)/8.00 =$	USD 6,595,254
upfront	$1600 \times 0.01/8.00 =$	2,000,000
total cost		8,595,254
equivalent annuity	8,595,254/5.551560664 =	1,548,259
same, in percent	1,548,259/200m =	% 0.774
`YIELD("1/1/2001","12/31/2007",0.0575,99,100,1,1)`	=	% 6.935
risk-free proxy		% 6.156
all-in spread		% 0.779

So far so good—but what is the point of the calculations? We do not really have a risk-free rate: 6.156 is a possible lower bound that is, however, rarely applied. If we had worked with 6.5 as the risk-free proxy, the equivalent annuity would have been calculated as 0.25% above 6.5%, and the IRR at 0.435% above it, not 0.77-0.78%. We could even make a case for using the midpoint rate rather than a borrowing rate. That midpoint rate would be 5.49, implying all-in spreads of 1.44%, as you can calculate.

DoItYourself problem 16.5

Do calculate the equivalent annuity part. (The IRR part is trivial, of course.)

As long as spreads are small and similar across countries, as in the USD-EUR case, these refinements hardly change the conclusions, but here we have a spread of 2.7 percent between bid and ask. We conclude that all calculations are, at best, tentative.

But lack of knowledge of the risk-free rate is not the only problem. Even if we had an active internal market for government bonds, the rate would still not be integrated with rates for other currencies, because worldwide financial investors cannot freely switch between CNY and USD lending (or borrowing, for that matter). The mechanism that normally equalizes the values and wipes out non-zero NPV's is missing, and along with that we lost all grounds to believe that the "risk-free" versions of the USD and CNY loans are truly equivalent. Remember that this last notion was the reason why only the PV'ed risk-spreads and costs need to be compared even for loans in different currencies. Conversely, without market integration the whole let's-just-compare-spreads approach is built on sand. Quicksand actually.

So all we can say is that, in terms of IRRs, USD borrowing would cost 4.84% while in CNY the figure is 6.94%. If the Yuan were expected to depreciate by about 2.0% per year, the two loans would be expected to have the same cost:

$$\text{break-even appreciation rate on FC, } a\text{:} \quad \underbrace{(1+a)(1+R^*)}_{\text{total return on FC}} = \underbrace{(1+R)}_{\text{return on HC}},$$

$$\Rightarrow a = \frac{1+R}{1+R^*} = \frac{1.0484}{1.0694} - 1 = -\%1.96,$$

In the case of the Yuan, at the time of writing the decision would be easy: the currency is undervalued by most standards; there is pressure from US Congress to revalue it, and the PBC seems to have chosen a course of slow and gentle appreciation. In short, the smart money would bet on an appreciation not depreciation of the Yuan against de US dollar. Given the extra 2% cost, there is no case for Yuan borrowing.

Of course, one is not always so lucky: with overvalued currencies (a policy often preferred by politicians in the past),[12] we'd have to weigh the cost of a high yield against the boon of expected depreciation. Signs of overvaluation would be a

[12] An overvalued home currency makes manufactured imports cheaper, which suits the city population and the political class; farmer exporters are hurt, but they often have less influence. An expensive currency rate also is regarded as adding prestige; devaluing would be an admission of defeat.

hamburger-parity rate far above those of comparable countries, or PPP rates that are unusually high; exchange controls; and interest-rate ceilings. But all this generates only hints and directions, not precise expected values. To make things worse, expectations are only part of the story: we should think of a normal risk premium too.

16.4 CFO's Summary

The main differences between international ("euro-") and domestic transactions are that the former are often extra-territorial, and the market is a liquid and unregulated wholesale market. As a result, spreads and costs are quite low, and the international markets have become an increasingly important source of funding for medium-size or large corporations. Apart from this, the transactions one can make in these markets are not fundamentally different from the transactions in standard domestic markets: there are time deposits and term loans, credit lines, and markets for bonds and short-term paper.

A more recent instrument is the forward or futures contract on interest rates, which we discussed in the Appendices to Chapters 4 and 6. Remember, from that discussion, that interest rates (spot and forward interest rates, and "yields at par" are all linked by arbitrage. Forward interest rates in various currencies are likewise linked through the forward markets.

Comparing loans is easy when markets are well developed and free. In that case, differences between risk-free rates should reflect the market's opinion about the currency, and switching between risk-free FC and HC lending would not affect value. The CFO's focus should therefore be on upfront costs and risk spreads. Using swaps, one can separate the currency of effective borrowing from the currency of effective exposure—for instance by borrowing at home and swapping into FC. So the rule is *always* to borrow where it is cheap in terms of costs and spreads, whether you fancy the currency or not. You can change the denomination afterwards via a swap, if you want. Hedging of operational exposure could be a consideration in the decision whether or not to swap the cheapest loan into another money. So could speculation—but bear in mind that the records of exchange-rate forecasters are patchy.

With less well-developed markets, the absence of a clear and market-set risk-free rate makes decisions much more difficult. One can compute total costs, but one often cannot separate out a risk-free component; and if a locally-set risk-free rate proxy is available after all, it still is unlikely to reflect a currency's relative prospects as viewed by the international market. If currency and risk-free interest rates reflect some officials' opinion rather than the market's views, the usual prior that financial deals are zero-NPV transactions would not even hold as a first approximation.

A sensible general prior might be that, for reasonably respectable companies, borrowing in developed markets should be more attractive, for the simple reason that sophisticated markets are cheaper to operate and its players better informed. Selectively subsidized loans in the host country could offset that, but the WTO frowns on practices like that. Interest rates that are capped without discrimination, in contrast, would be acceptable to the WTO, and still exist in some places. Another item that could tilt the balance back to the host-country market is the exchange rate. Controlled exchange rates often imply one-way bets: it is usually obvious whether the currency is overvalued or the converse. But remember that getting the timing and size of the adjustment right remains difficult. There is no easy way out, here. For decisions like this, CFOs will not be replaced by computers any time soon.

16.5 Test Your Understanding

16.5.1 Quiz Questions

The questions also cover interest forwards and futures, which were discussed in the appendix to Chapter 4.

True-False Questions

1. The abolition of the Interest Equalization Tax, Regulation Q, the cold war, and the US and UK foreign exchange controls have taken away most of the reasons why euromarkets exist. As a result, we can expect these markets to decline in the near future.
2. Without the US trade deficit, the euromarkets would have developed more slowly.
3. With a floating-rate loan, the bank is free to adjust the interest rate at every reset date in light of the customer's creditworthiness.
4. One of the tasks of the lead bank under a syndicated bank loan is to make a market, at least initially.
5. The purpose of using a paying agent is to reduce exchange risk.
6. Caps and floors are options on interest rates. Because interest rates are not prices of assets, one cannot price caps and floors using an option pricing model that is based on asset prices.
7. Because euroloans are unsecured, the spread over the risk-free rate is a very reliable indicator of the borrower's general creditworthiness.
8. FRAs are not really a good hedge against future interest rates because one does not actually make the deposit or take up the loan.
9. A note-issuing facility forces the borrowing company to borrow at a constant spread, while a revolving underwritten facility gives the borrower the benefit of decreasing spreads without the risk of increasing spreads.
10. The fact that eurobonds are bearer securities makes them less attractive to most investors.
11. Bond stripping is always done with a pair of scissors: you just clip off the coupons.
12. Disintermediation is the cause of the lower creditworthiness of banks, and has lead to capital adequacy rules.
13. Ignoring the small effects of marking to market, the standard quote for a eurocurrency futures price is basically a forward price on a CD.

Multiple-Choice Questions

1. Eurocurrency and euroloan markets are attractive because:
   (a) the spread between the buy and ask exchange rates is lower than in the interbank exchange market.
   (b) the bid-ask spread between the lending and borrowing interest rates is lower.
   (c) eurobanks are not subject to reserve requirements.
   (d) eurobanks are not subject to capital adequacy rules (the so-called BIS rules).

2. Eurobanks borrow for short maturities and lend for longer maturities. They can reduce the interest risk by:
   (a) extending fixed-rate loans.
   (b) extending floating-rate loans.
   (c) extending revolving loans.
   (d) shorting forward forwards (that is, getting a forward contract on a loan, not on a deposit).
   (e) shorting in FRAs.
   (f) going long eurocurrency futures.
   (g) buying forward the currency in question.

3. A cap on a floating-rate euroloan:
   (a) protects the borrower against high short-term interest rates.
   (b) protects the lender against high short-term interest rates on the funding side.
   (c) is similar to a call option on short-term paper with the cap rate, as nominal rate; and the borrower is the holder of the call option.
   (d) is similar to a put option on short-term paper with the cap rate, as nominal rate; and the borrower is the holder of the put option.
   (e) is similar to a put option on short-term paper with the cap rate, as nominal rate; and the lender is the holder of the put option.

4. Which of each pair best describes eurobanking?
   (a) retail/wholesale
   (b) individual lender/bank consortium
   (c) reserve requirements/limited or no reserve requirements
   (d) unsecured/secured
   (e) fixed-rate lending/floating-rate lending

(f) foreign exchange markets/money markets

(g) open to all companies/open to the better companies only

5. Matching Questions: Choose from the following list of terms to complete the sentences
below: paying agent, managing banks, trustee bank, placing agents, market, lead bank (or lead manager), participating banks, prospectus, gray market, fiscal agent, buy forward, underwrite, lead manager, red herring.

   A consortium (or syndicate) that extends a euroloan consists of many banks that could play different functions. In a euroloan, the (a) negotiates with the borrower for tentative terms and conditions, obtains a mandate, and looks for banks to provide the money or undertake to provide the money if there is any shortfall in funds. The banks that provide the actual funding are called (b). Because at the time of the negotiations the funding is not yet arranged, the (c) often contacts a smaller number of (d) banks who (e) the loan, that is, guarantee to make up for the shortage of funds if there is any such shortfall. The (f), finally, is the bank that receives the service payments from the borrower and distributes them to the participating banks.

   Placement of eurobonds is most often via a syndicate of banks or security houses. The lead bank or (g) negotiates with the borrower, brings the syndicate together, makes a (h) (at least initially), and supports the price during and immediately after the selling period. There are often, but not always, (i) that underwrite the issue and often buy part of the bonds for their own account. The (j) call their clients (institutional investors or individuals) and sell the bonds on a commission basis. The (k) takes care of withholding taxes, while the (l) monitors the bond contract. Prospective customers can find information about the issuing company and about the terms and conditions of the bond in the (m). Often an unofficial version of the prospectus is already circulating before the actual prospectus is officially approved; this preliminary prospectus is called the (n). On the basis of this document, investors can already (o) the bonds for a few weeks before the actual issuing period starts. This period of unofficial trading is called the (p) period.

16.5.2 Applications

1. You are an A-quality borrower, and you pay 10 percent on a five-year loan with one final amortization at the end and annual coupons. This is 1 percent above the spread paid by an AAA borrower. What will be the up-front fee for which your bank should be willing to lower the rate by 1 percent?

2. A bank offers you the following rates for a 5-year loan with annual coupons: 10 percent fixed, or (when you borrow floating-rate) LIBOR + 2 percent. You prefer to borrow floating-rate, as you expect a drop in interest rates. Another bank offers you LIBOR + 1.5 percent, but asks a substantial up-front fee. How can you compute which bank offers the better terms?

3. You bought an option that limits the interest rate on a future six-month loan to, at most, 10 percent *p.a.*

   (a) If, at the beginning of the six-month period, the interest rate is 11 percent, what is the expiration value of this option?

   (b) What is the option's expiration value if the interest rate turns out to be 8 percent?

4. You bought an option that limits the interest rate on a future six-month deposit to at least 10 percent *p.a.*

   (a) If, at the beginning of the six-month period, the interest rate is 11 percent, what is the market value of this option?

   (b) What is the option's value if the interest rate turns out to be 8 percent?

Chapter 19

Setting the Cost of International Capital

This chapter deals with how to set the cost of capital, which is the discount rate used in Capital Budgeting. The chapter title adds one word: international. Note that what is said to be international is *capital*. The title does not say "the international cost of capital", as the 1994 book (and many others) did: such a title would have suggested that there is something like a national cost (for domestic projects, presumably) and, next to that, an international cost, for transborder investments. No, there is just one cost of capital, and that capital is international. Shares of large corporations are held by people everywhere; and, equally important, even shareholders of smaller, more locally held firms still invest part of their wealth in foreign stocks. This has two implications for the way the cost of capital is to be set. First, managers should ask the question how much risk this project adds to an internationally diversified portfolio instead of to a local market portfolio (the traditional method), and set a cost of capital that is commensurate with this international risk. Second, management has to take into account that the expected return differs depending on what currency the shareholder uses as the (quasi-)real numeraire; and so does the risk-free rate that serves as one benchmark item in the model. That is, when setting the cost of capital, the issue of exchange risk has to be taken into account too, in the sense of investors having different numeraires in which they are doing, or supposed to be doing, their optimum-portfolio calculations.

There is a second—and largely independent—issue related to exchange rates: how do we bring expected cash flows and cost of capital in line with each other. The issue arises because when, say, an Australian firm invests in India, the expected future cash flows are normally first expressed in Rupees. Yet, the argument typically goes, the Australian owners care about Australian Dollars only—we'll make this argument more precise as we proceed—and the cost of capital we would estimate is probably expressed in AUD. One cannot discount INR cash flows using an AUD discount rate. So at one point we need to go from INR to AUD.

There seems to be two ways we could go about this, similar to what we did earlier for risk-free cash flows. As we saw, a risk-free claim on INR 1 can be PV'ed in INR terms first, by discounting the INR cash flow (unity) at the INR risk-free rate and this value is then translated into AUD at the going spot rate. Alternatively, we can translate the future cash flow into AUD using the expected future spot rate, and then discount at an AUD rate that takes into account the risk. Both are linked via the forward rate as the risk-adjusted expectation and CIP:

$$\begin{aligned} \frac{E_t(\tilde{S}_T)}{1+r_{t,T}+RP_{t,T}} &= \frac{F_{t,T}}{1+r_{t,T}}, && \text{(F=CEQ)} \\ &= \frac{1}{1+r^*_{t,T}} S_t. && \text{(CIP)} \qquad (19.1) \end{aligned}$$

Similarly then, in case of a risky FC cash flow, we could first translate the *future* INR cash flows into AUD using the expected future spot rate, and then PV these using an AUD discount rate, set *e.g.* on the basis of the standard Capital Asset Pricing Model (CAPM), the way Australians would value a domestic Australian project. Alternatively, we could argue that the Australian ownership hardly matters, and simply conduct the entire cost-benefit analysis in INR, the way an Indian owner would do: take INR cash flows, and discount at the rupee rate of return. Having found the value in INR, we then translate the *present* value into AUD. And if that second solution really works, exchange-rate forecasts and currency risk can be totally eliminated from the analysis, it would seem.

In this chapter we show that the above analysis is quite incomplete. The main lessons to be remembered from this chapter are the following:

- **Translation of FC cash flows requires more than just an expected exchange rate** Suppose we follow the first route and translate our Rupee cash flows, $\tilde{C}^*_T$, into AUD. What we need are expected AUD cashflows; but the expectation of a product, $\mathrm{E}(\tilde{C}^*_T\tilde{S}_T)$ involves not just the expectations of $\tilde{C}^*_T$ and $\tilde{S}_T$, but also the covariance between the two.

  This, at first sight, makes the first route even more difficult. All the more reason to go for the alternative one, then? Unfortunately, this alternative would not always work:

- **Host-currency v home-currency valuation** Valuation in Rupees, the way an Indian investor would do it—using the Rupee risk-free rate and a premium for market risk measured in Rupees—should produce the same result, after translation, as valuation *à l'Australienne* only if the Indian and Australian capital markets are well integrated. Indeed, if investors from each country can freely invest in each other's market (and possibly in other markets), arbitrage flows would occur if the value to Australians were different from the value to Indian investors (after translation into a common currency).

  In the case of India integration of the capital market into the mainstream work market is doubtful, for the time being. But even if it were true, the Indian-

Rupee approach would still not exonerate us from thinking about expected exchange rate changes and exchange covariance risk:

- **In open markets, exchange risk affects any cost of capital ...** In principle, exchange risk enters asset pricing as soon as the investor base for which we want to value the project is part of an international market. Thus, Australia being part of a nearly worldwide capital market, an International CAPM (i-CAPM) should be used whether the project is situated at home or abroad. Intuitively, in an international capital market, asset prices result from the interaction of portfolio decisions by people from many different countries, each having their own currency. Exchange risk makes people disagree about expected returns and risk; for example, the AUD treasury bill is risk-free to Australians, but not to Canadians or Japanese. This heterogeneity of perspectives does affect asset pricing, and introduces currency risk premia into the CAPM, in principle one for each currency area that is part of the international capital market.

  Thus, in a way things are even more complicated than your worst fears might have been: you need expected returns on *all* currencies in the international capital market, and covariances for your project with each of these currencies. In addition, exposures to exchange rates are even harder to estimate than betas. Fortunately, ...

- **... but currency risk premia are small** The literature on the forward rate as a predictor of futures spot rates shows that, while the currency risk premium is surely not a constant, it is small and seems to fluctuate around zero. So one could use a shortcut, omitting the forex items in the i-CAPM formula, so that it looks rather like the familiar domestic CAPM. Two differences remain: the market portfolio is the world-market index rather than a domestic one, and the market beta is from a multiple regression with all exchange rates included.

The discussion can be summed up as follows.

(*i*) **Which CAPM** You use a (possible simplified) i-CAPM when the *home* country is part of an international capital market; the domestic model works only for segmented home markets.

Note, incidentally, that this holds for any investment, whether at home or abroad.

(*ii*) **Which currency** If home and host are *both* part of the same international market, either currency will do for valuation purposes; otherwise only the investors' HC can be used.

This chapter addresses these issues in the following order. First we discuss the effect of capital-market integration or segmentation on the capital-budgeting procedure (Section 19.1), notably which should come first, translation from FC to HC or discounting. The bulk of the chapter then relates to the determination of the cost of capital. In the second section, we present the traditional single-country CAPM, starting from the efficient-portfolio problem familiar from basic finance courses. In Section 19.3, we explain how to modify this model when assets are priced in an international market. The case that we discuss is one where capital markets are integrated across many countries, but where imperfections in the goods markets create real exchange risk. Section 19.4 concludes with a review of the implications of this chapter for capital budgeting.

19.1 The Link between Capital-market Segmentation and the Sequencing of Discounting and Translation

To initiate our discussion of the effect of capital market integration or segmentation on the capital budgeting procedure, we explain why capital budgeting can be done in terms of foreign currency when the home- and host-country capital markets are integrated, and how the procedure is to be modified when the home- and host-country capital markets are segmented from each other.

Almost inevitably, capital budgeting starts with cash-flow projections expressed in host (foreign) currency. When one prepares cash flow forecasts there is no real choice but to start from currently prevailing prices for similar products in foreign currency. On the basis of this you set your own price(s), taking into account the positioning of the product(s). Then you try to figure out production costs on the basis of data from similar plants and local wages and other input costs. (Don't forget the initial inefficiencies, the learning curve. And think of possible price drops later when competition catches up or the rich segment has been creamed off or excitement about your product wanes.) This way you obtain cash-flow forecasts, all typically at current (i.e. constant) FC prices. Finally you adjust the figures for expected foreign inflation. This practice stems from the empirical fact, noted in Chapter 3, that prices in any given country are sticky (apart from general inflation) and to a large extent independent of exchange rate changes.

DoItYourself problem 19.1

You could think of an alternative version of the final step: translate the constant-prices cash flow into HC and then adjust for inflation in the investor's home country. Show that this unattractively assumes relative PPP, at least as an expectation. Assume risk-free cash flows at constant FC prices, for simplicity.

So we usually end op with expected cash flows in FC. However, the ultimate purpose of capital budgeting is to find out whether the project is valuable to the parent company's shareholders. The correct procedure is to see how *they* price similar existing projects. We can see that only by looking at their own capital market; that is, we use the shareholders' home capital market to get the risk-free rate and the estimated risk premium. But this delivers a cost of capital in HC units, which can only be used to discount HC expected future cash flows. For example, one would not use a low JPY-based discount rate to PV a stream of Zimbabwe Dollar cash flows. In short, although the natural input data are cash flow forecasts expressed in foreign currency, in principle we have to make the translation from foreign currency to home currency before we can discount. To what extent would it be acceptable, instead, to discount FC cash flows at a FC rate, and then to translate the FCPV into HC using just the current spot rate? After all, this is the way a local investor goes about the valuation.

This type of valuation in foreign currency, as if the owner were a host-country investor, is correct if the host- and home-country financial markets are integrated, that is, if there are no restrictions on cross-border portfolio investment between the two countries and if investors effectively hold many foreign assets. Indeed. the implication of market integration is that all investors, regardless of their place of residence, use the same cost of capital when they compute the price of any given asset (in some given common currency) from the expected cash flows of this asset (expressed in the same common currency). One way to explain this claim is by contradiction. If investors from countries A and B used a different cost of capital when computing the price of some given asset (in some given base currency) from the asset's expected cash flows (measured in the same base currency), then the price of the asset in country A would differ from the price of the same asset in country B. The resulting arbitrage opportunities would lead to international trading in the shares until the price difference disappeared. By equating prices across countries, international arbitrage also equates the costs of capital that various investors use when linking the asset's price to the expected cash flows paid out by the asset.[1] Thus, in integrated markets, a home-country investor and a host-country investor fully agree about the project's value.

In the perfect-markets approach of Chapter 4, perfected integration was taken for granted. But in the case of FDI into emerging countries it is not always obvious that integration is a reasonable approximation, even though restrictions are gradually being abolished in many countries. The problem is that in segmented markets one cannot simply value a foreign cash flow as if it were owned by host-country investors. In the absence of free capital movements, there is no mechanism that equates prices and discount rates across the two markets. Thus, to the managers of the parent firm, the relevant question becomes: What price would home-country investors normally be prepared to pay for the project? As we saw, the way to proceed is to identify cash flow patterns that have similar risks and that are already priced in the home-country capital market. Once we have identified a similar asset that is already priced in the home capital market, we can then use the same discount rate for the project that

[1]Investors that are not willing to pay a high price then sell to others that are. Portfolio rebalancing also modifies the risk: the risk of holding Samsung shares is very different depending whether this company represents 90 % of one's portfolio versus just 0.1% of a well-diversified package of securities. So reducing the weight of one asset, and replacing it by others that offer more diversification, lowers required returns for that asset and increases the price one is willing to pay for it. In the end, when both domestic and foreign investors hold very similar portfolios, required returns would converge.

we want to value as that for the traded assets. To implement this procedure, we need a theory, like the Capital Asset Pricing Model, to tell us what types of risk are relevant, how these risks should be measured, and what return is expected in the home-country capital market in light of the project's risks. Since we use the home-country capital market as the yardstick, the discount rate is the required return in home currency—and if the cost of capital is expressed in home currency, we have to translate the expected cash flows and their risks from foreign currency into home currency before we discount.

For such a translation, we need expected values for the future spot rates for various maturities. In fact, we need also the covariance. If $\tilde{C}^*$ denotes the cash flow in FC and $\tilde{C}^*\tilde{S}$ the cash flow in HC, then

$$\mathrm{E}(\tilde{C}^*\tilde{S}) = \mathrm{E}(\tilde{C}^*) \times \mathrm{E}(\tilde{S}) + \mathrm{cov}(\tilde{C}^*, \tilde{S}). \tag{19.2}$$

You may have noticed the covariance effect in the Freedonian Crown exposure example in Chapter 13, which we reproduce here:

Example 19.1

A British company is considering a project in Freedonia. Assume that the Freedonian crown (FDK) cash flow can take on either of two equally probable values, FDK 150 or FDK 100, depending on whether the Freedonian economy is booming or in a funk. Let there also be two, equally probable time-T spot rates, GBP/FDK 1.2 and 0.8. Thus, measured in GBP, there are four possible cash flows: 150 × 1.2 = GBP 180, 150 × 0.8 = GBP 120, 100 × 1.2 = GBP 120, and 100 × 0.8 = GBP 80. These numbers are shown in Table 19.1. In each cell, we also show the joint probability of each particular combination. When the FDK is expensive, a recession is more probable than a boom because an expensive currency means that Freedonia is not very competitive. The inverse happens when the crown is trading at a low level; then it is more likely that the economy will be booming. These effects are reflected in the probabilities shown in each of the four cells in Table 19.1.

Table 19.1: **Cash flows for the Freedonian project**

	State of the economy			
	Boom: C^*=150	Slump: C^*=100	Prob(S)	E($\tilde{C}\vert S$)
S_T=1.2	p = 0.15; C=180	p=0.35; C=120	0.50	138
S_T=0.8	p = 0.35; C=120	p= 0.15; C= 80	0.50	108
Prob(C^*)	p = 0.50	p = 0.50		

The expectations of the exchange rate and the FDK cash flows are easily calculated as

$$\mathrm{E}(\tilde{S}) = (0.50 \times 1.2) + (0.50 \times 0.8) = 1.00, \tag{19.3}$$

$$\mathrm{E}(\tilde{C}^*) = (0.50 \times 150) + (0.50 \times 100) = 125. \tag{19.4}$$

But the expected cash flow is not 1.00 × 125 = 125:

$$\mathrm{E}(\tilde{S}\tilde{C}^*) = (0.15 \times 180) + (0.35 \times 120) + (0.35 \times 120) + (0.15 \times 80) = 123. \tag{19.5}$$

The shortfall of 2 (=125 − 123) is due to the fact that high FDK cash flows tend to go together with low exchange rates and vice versa. This effect is lost if one just multiplies through the two expectations, because that computation implicitly assigns probabilities 0.25 to each cell:

$$\begin{aligned} \mathrm{E}(\tilde{S})\mathrm{E}(\tilde{C}^*) &= [(0.50 \times 1.2) + (0.50 \times 0.8)] \times [(0.50 \times 150) + (0.50 \times 100)], \\ &= (0.25 \times 180) + (0.25 \times 120) + (0.25 \times 120) + (0.25 \times 80). \end{aligned} \tag{19.6}$$

So when we use the Translate First approach, the expected GBP cash flow is GBP 123 not 125.[2] This number is to be discounted at the appropriate home currency discount rate, that is, the GBP risk-free rate plus a risk premium that reflects the risk of the GBP cash flows to the British investor.[3] The Capital Asset Pricing Model, to be discussed in Sections 19.2 and 19.3, provides a way to estimate the appropriate discount rate.

While the Translate First approach is very general, it requires explicit exchange-rate forecasts, and the covariance. These do not come in explicitly if we take the Discount First route, and compute a PV for the expected flow $\mathrm{E}(\tilde{C}^*) = 125$, using the FDK risk free rate and risk premium. This would be all right if the Freedonian and British markets are well integrated.

We have seen how to obtain expected cash flows, but not how to obtain appropriate discount rates when cash flows are risky. This is the task in the remainder of this chapter. Section 2 reviews the single-country CAPM. Section 3 extends the model to a multi-country setting.

19.2 The Single-Country CAPM

Our discussion of the traditional (single-country) CAPM starts from asset demand theory. The key assumption of this asset demand theory is that investors rank portfolios on the basis of two numbers, the expected nominal portfolio return and the variance of the nominal portfolio return. Implicit in the use of nominal returns is an assumption that inflation is deterministic, or at least that inflation uncertainty has little impact on asset pricing. The theory of optimal portfolios, as developed by Markowitz (1952), can also be interpreted as a theory that tells us how expected

[2] In the above example, the cov-correction is relatively small. But the link between exchange rate and cash-flow is weak too, in the above story: it just works via general economic activity. In reality, there often is a strong, direct link, for instance if the firm is an exporter or importer, and then the covariance would be bigger.

[3] Recall that if capital markets within, say, the OECD are well integrated, the UK value would also be correct for any other investor from any other OECD country. (The OECD is just a for-example term: the world market now counts many non-OECD members.)

returns are related to risk in an efficient portfolio. This relationship is due to Sharpe (1964), Lintner (1965), and Mossin (1965).

19.2.1 How Asset Returns Determine the Portfolio Return

The model is typically derived in terms of returns rather than prices: academics use returns in empirical work, and practitioners want a formula for the expected return to be used for NPV applications. The key relation is that the realized return on the portfolio (subscript p) can always be written as (i) the risk-free return over that period, plus (ii), for all risky assets in the portfolio, the weighted average of the returns over and above the risk-free rate:

$$\tilde{r}_p - r = \sum_{j=1}^{N} x_j(\tilde{r}_j - r), \tag{19.7}$$

with a weight x_j defined as the initial amount invested in asset j, divided by total initial investment. A return over the risk-free rate is called an *excess return*, and its expected value is called the *risk premium.*

Example 19.2

You have 1000 to invest. Below, we show for three risky assets (denoted as 1, 2, 3) an initial price, the number of shares you buy, your total initial investment per asset, the asset weight, a possible time-1 price, the corresponding return, and the weighted return. The risky assets take up 900 of the money, so the balance, 100, is invested risk-free at, say, 5 percent. In the table we see the weights,[4] and how they sum to unity. We next compute the value of the portfolio at time 1, and see that it has gone up to 1105, implying a (net rate of) return of 0.105, *i.e.* 10.5 percent. The excess return is $10.5 - 5 = 5.5\%$, and this is exactly what you get by summing the value-weighted "excess" returns on the three risky assets.

	time-0 data and decisions					time-1 result		(excess) rates of return		
	j	$V_{j,0}$	n_j	$n_jV_{j,0}$	x_j	$V_{j,1}$	$n_jV_{j,1}$	r_j	$r_j - r$	$x_j(\tilde{r}_j - r)$
risky:	1	100	4	400	0.40	120	480	0.20	0.15	0.060
	2	50	4	200	0.20	70	280	0.40	0.35	0.070
	3	25	12	300	0.30	20	240	-0.20	-0.25	-0.075
subtotal				=900	=0.90					=0.055
risk-free	0			+100	+0.10		105			+0.05
total				=1000	=1.00		=1105			r_p=0.105

[4]Note that the weights we need in the formula are initial weights, determined by time-0 numbers, meaning that they are not stochastic.

Figure 19.1: **Combinations of risky stock portfolio s and asset 0**

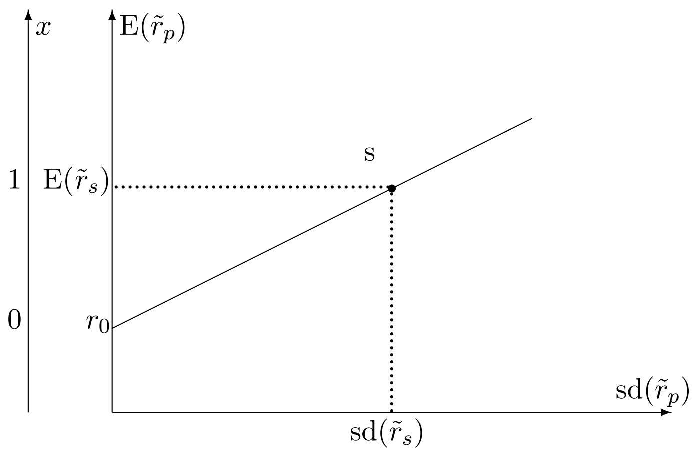

DoItYourself problem 19.2

Rework the example by changing the initial investment in asset 1 from 400 to 800, maintaining the other risky positions and adjusting the risk-free one. Check that the weighted excess-return formula still gives the right answer.

19.2.2 The Tangency Solution: Graphical Discussion

Consider the feasible combinations of expected return and standard deviation. The simplest case is one with a risk-free asset, subscripted "0", and a risky stock denoted as s. We invest a fraction x into the risky stock portfolio with return $\tilde{r}_s$ while $1-x$ is invested risk-free. The portfolio return is

$$\tilde{r}_p = x\tilde{r}_s + (1-x)r_0 = r_0 + x\,(\tilde{r}_s - r_0) \Rightarrow \begin{cases} \mathrm{E}(\tilde{r}_p) = r_0 + x\,\mathrm{E}(\tilde{r}_s - r_0), \\ \\ \mathrm{sd}(\tilde{r}_p) = |x|\,\mathrm{sd}(\tilde{r}_s) \end{cases} \tag{19.8}$$

So for non-negative values of x, both expected return and standard deviation are linear functions of x. This will imply that all (E, sd) combinations achievable with the risk-free asset and the risky portfolio are found on a halfline. To show this we use a trick that is often applied in thermometers, where heat is measured on two scales, say Celsius and Fahrenheit, that are linearly related. Same here: x and $\mathrm{E}(\tilde{r}_p)$ are linearly related, so we can show them as two scales on one axis, as we do in Figure 19.1. To link the x and E_p scales we calibrate them using any two known corresponding points: $x = 1$ means $\mathrm{E}(\tilde{r}_p) = \mathrm{E}(\tilde{r}_s)$ while $x = 0$ means $\mathrm{E}(\tilde{r}_p) = r_0$. All this gives us the double-scaled axis shown in Figure 19.1. If $\mathrm{sd}(\tilde{r}_p)$ is linear in x_+, then looking at the other scale of the axis we must conclude it is linear in $\mathrm{E}(\tilde{r}_p)_{>r_0}$ too. The sd values for the calibration points are 0 and $\mathrm{sd}(\tilde{r}_s)$, respectively, and all risk-return combinations for intermediate or higher values of x or $\mathrm{E}(\tilde{r}_p)$ are on one

Figure 19.2: **The risk-return bound with just risky assets**

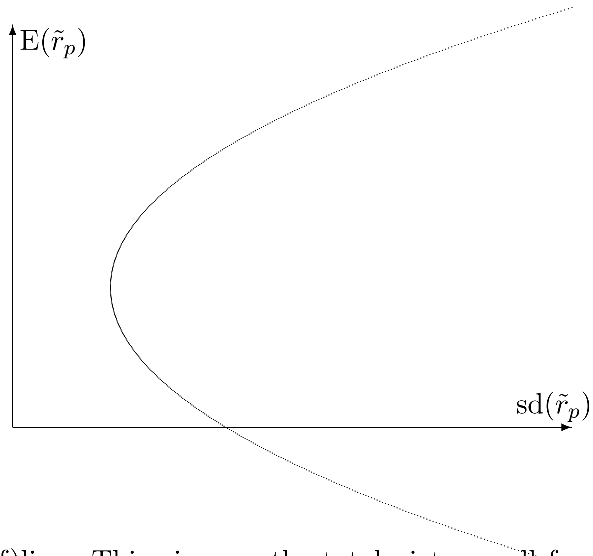

and the same (half)line. This gives us the total picture: all feasible combinations with $x \geq 0$ are on a halfline from $(\mathrm{sd}(\tilde{r}_p), \mathrm{E}(\tilde{r}_p)) = (0, r_0)$ through $(\mathrm{sd}(\tilde{r}_s), \mathrm{E}(\tilde{r}_s))$. The slope of that halfline is called the *Sharpe Ratio*:

$$\forall x \geq 0 : \ \frac{\mathrm{E}(\tilde{r}_p - r_0)}{\mathrm{sd}(\tilde{r}_p)} = \frac{\mathrm{E}(\tilde{r}_s - r_0)}{\mathrm{sd}(\tilde{r}_s)} = s\text{'s Sharpe Ratio.} \tag{19.9}$$

Now look at a second simple case, where the portfolio consists of two imperfectly correlated risky assets, subscripted 1 and 2. Now we have

$$\tilde{r}_p = x_1\tilde{r}_1 + (1 - x_1)\tilde{r}_2; \tag{19.10}$$

$$\Rightarrow \quad \left\{ \begin{array}{l} \mathrm{E}(\tilde{r}_p) = \mathrm{E}(\tilde{r}_1) + x_1\,[\mathrm{E}(\tilde{r}_1) - \mathrm{E}(\tilde{r}_2)], \\ \\ \mathrm{sd}(\tilde{r}_p) = \sqrt{x_1^2 \mathrm{var}(\tilde{r}_1) + 2x_1(1 - x_1)\mathrm{cov}(\tilde{r}_1, \tilde{r}_2) + (1 - x_1)^2 \mathrm{var}(\tilde{r}_2)} \end{array} \right. \tag{19.11}$$

From the first implication we conclude that x_1 and expected return are still two sides of the same thermometer. The sd function looks messier. But we immediately see that variance is quadratic in x_1 and, therefore, in $\mathrm{E}(\tilde{r}_p)$ too. This means a rotated "U"-shape-like graph (or a rounded V, if you want) opening towards the right. Warping the risk axis by taking squareroots does not fundamentally change the shape of the relation, as you can check using a spreadsheet. We end up with a feasible set like in Figure 19.2. Basic textbooks will tell you that, if there are more than two risky assets, the feasible combinations in a (std, E) space graph is still similar.

The last step is to look at N risky assets and a risk-free one. We return to Figure 19.1 except that the risky part of the portfolio, s, must be chosen from a feasible set shaped like in Figure 19.2. A risk-averse mean-variance investor wants

Figure 19.3: **Efficient Portfolios & the Tangency Portfolio**

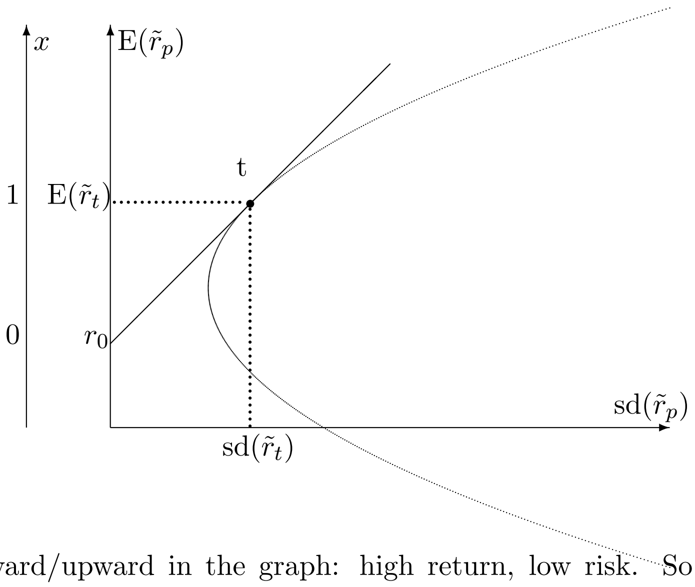

to be leftward/upward in the graph: high return, low risk. So s will be chosen from the left-upper risky-asset bound. Among all such portfolios, the best one is the portfolio that rotates the halfline from $(\text{sd} = 0, \text{E} = r_0)$ as far upward/leftward as is feasible—the one with the highest Sharpe Ratio. It follows that the optimal choice is the *tangency portfolio*, the one where the halfline from $(\text{sd} = 0, \text{E} = r_0)$ just touches the V-curve that bounds the risky-assets risk-return set. All portfolios on this halfline are efficient. They all are combinations of the risk-free asset and the *tangency portfolio*, subscripted t.

We now want to take a peek at the analytical solution and its implication. To understand how the tangency portfolio can be found we need to understand first how a small change in one of the portfolio weights affects the expected return and the variance of the portfolio return.

19.2.3 How Portfolio Choice Affects Mean and Variance of the Portfolio Return

We want to understand what happens if investors choose portfolios on the basis of the mean and variance of the portfolio return. To figure out how these people think, we need to understand how portfolio choice affects the mean and variance of the total return. The link is, of course, Equation [19.7]: $\tilde{r}_p = r + \sum_{j=1}^{N} x_j\,(\tilde{r}_j - r)$. From this it follows that

$$\text{E}(\tilde{r}_p) = r + \sum_{j=1}^{N} x_j\,\text{E}(\tilde{r}_j - r), \tag{19.12}$$

$$\text{var}(\tilde{r}_p) = \sum_{j=1}^{N} x_j \sum_{k=1}^{N} x_k\,\text{cov}(\tilde{r}_j,\, \tilde{r}_k). \tag{19.13}$$

The first formula is pretty obvious. To interpret the second one, it helps to derive it in two steps, as follows:[5]

$$\mathrm{var}(\tilde r_p) = \mathrm{cov}(\tilde r_p,\, \tilde r_p) = \mathrm{cov}(\sum_{j=1}^{N} x_j \tilde r_j,\, \tilde r_p) = \sum_{j=1}^{N} x_j\, \mathrm{cov}(\tilde r_j,\, \tilde r_p), \tag{19.14}$$

where

$$\mathrm{cov}(\tilde r_j,\, \tilde r_p) = \mathrm{cov}(\tilde r_j,\, \sum_{k=1}^{N} x_k \tilde r_k) = \sum_{k=1}^{N} x_k\, \mathrm{cov}(\tilde r_j,\, \tilde r_k). \tag{19.15}$$

This tells you that the portfolio variance is a weighted average of each asset's covariance with the entire portfolio; and each of these assets' portfolio covariances is, in turn, a weighted average of the asset's covariance with all components of the portfolio.

Example 19.3

We compute the portfolio expected excess return, the assets' covariances with the portfolio return, and the portfolio variance when the risky assets' weights are 0.50 and 0.40 (implying $x_0 = 0.10$):

	$\mathrm{E}(\tilde r_j - r)$	$\mathrm{cov}(\tilde r_j, \tilde r_1)$	$\mathrm{cov}(\tilde r_j, \tilde r_2)$
1	0.200	0.16	0.05
2	0.122	0.05	0.09

$$\begin{aligned}
\mathrm{E}(\tilde r_p - r) &= 0.50 \times 0.200 + 0.40 \times 0.122 = 0.1488 \\
\mathrm{cov}(\tilde r_1, \tilde r_p) &= 0.50 \times 0.160 + 0.40 \times 0.050 = 0.1000 \\
\mathrm{cov}(\tilde r_2, \tilde r_p) &= 0.50 \times 0.050 + 0.40 \times 0.090 = 0.0610 \\
\Rightarrow \mathrm{cov}(\tilde r_p, \tilde r_p) &= 0.50 \times 0.100 + 0.40 \times 0.061 = 0.0744
\end{aligned}$$

How do these numbers change when the portfolio weights are being tweaked? First look at a two-risky-assets example and see how mean and variance are affected by a small change in the weight of asset 1 (implicitly matched by a small offsetting change in the weight for the risk-free bond, asset zero):

$$\begin{aligned}
\mathrm{E}(\tilde r_p - r) &= x_1\, \mathrm{E}(\tilde r_1 - r) + x_2\, \mathrm{E}(\tilde r_2 - r); \\
\Rightarrow \frac{\partial \mathrm{E}(\tilde r_p - r)}{\partial x_1} &= \mathrm{E}(\tilde r_1 - r),
\end{aligned} \tag{19.16}$$

[5] We use the fact that, inside a variance, risk-free returns added or subtracted play no role: $\mathrm{var}(r + \sum x_j(\tilde r_j - r)) = \mathrm{var}(\sum x_j \tilde r_j)$.

and

$$\begin{aligned}
\text{var}(\tilde{r}_p) &= x_1^2\text{var}(\tilde{r}_1) + 2x_1x_2\text{cov}(\tilde{r}_1,\tilde{r}_2) + x_2^2\text{var}(\tilde{r}_2);\\
\Rightarrow \frac{\partial\text{var}(\tilde{r}_p)}{\partial x_1} &= 2x_1\text{var}(\tilde{r}_1) + 2x_2\text{cov}(\tilde{r}_1,\tilde{r}_2),\\
&= 2[x_1\text{cov}(\tilde{r}_1,\tilde{r}_1) + x_2\text{cov}(\tilde{r}_1,\tilde{r}_2)],\\
&= 2[\text{cov}(\tilde{r}_1,x_1\tilde{r}_1) + \text{cov}(\tilde{r}_1,x_2\tilde{r}_2)],\\
&= 2\text{cov}(\tilde{r}_1,x_1\tilde{r}_1 + x_2\tilde{r}_2),\\
&= 2\text{cov}(\tilde{r}_1,\tilde{r}_p). \qquad (19.17)
\end{aligned}$$

Similarly, $\frac{\partial \text{E}(\tilde{r}_p - r_0)}{\partial x_2} = \text{E}(\tilde{r}_2 - r)$ and $\frac{\partial\text{var}(\tilde{r}_p)}{\partial x_2} = 2\text{cov}(\tilde{r}_2,\tilde{r}_p)$.

DoItYourself problem 19.3

Recompute the expected excess return and variance when, in the previous example, x_1 is increased from 0.50 to 0.51. Check how the scaled change in the mean, $\Delta\text{E}/\Delta x_1$, is exactly the first asset's own expected excess return. Likewise, check how the scaled change in the variance, $\Delta\text{var}/\Delta x_1$, is about twice the first asset's own covariance with the portfolio return.[6]

DoItYourself problem 19.4

Consider a portfolio with, initially, $x_1 = 0.5$ and $x_2 = 0$ so that $\text{var}(\tilde{r}_p) = \text{var}(0.5\,\tilde{r}_1) = 0.5^2\,\text{var}(\tilde{r}_1)$. Then increase the second weight to 0.001. Write out the change in the variance, and check whether it is far from $2\text{cov}(\tilde{r}_2,\tilde{r}_p) \times 0.001$.

19.2.4 Efficient Portfolios: A Review

Recall that a portfolio is efficient if it has the highest expected return among all conceivable portfolios with the same variance of return. We just reviewed the probably familiar result that any efficient portfolio is a combination of two building blocks: the risk-free asset, and the tangency portfolio of risky assets (Figure 19.3). But what is perhaps less obvious is how the tangency portfolio must be constructed and what this implies for the risk-return relation. Let us consider this.

It is easily shown that, if a portfolio is to be efficient, then for each and every asset the marginal risk-return ratio—the ratio of any asset's marginal "good" (its contribution to the portfolio's expected excess return) to the asset's marginal "bad"

[6] For the variance, the scaled difference is not perfectly the same as the partial derivative because the function is quadratic in the weights, not linear. (For non-linear functions, obviously, $\Delta y/\Delta x \neq \text{d}x/\text{d}y$.) But note how the scaled change in fact equals the average of the original and the revised covariances (0.1000 when $x_1 = 0.50$, and 0.1016 when $x_1 = 0.51$). In the limit, the two covariances are so close that they become indistinguishible from their average.

(its contribution to the portfolio's risk)—must be the same, see TekNote 19.1. We just identified the asset's contribution to the portfolio's expected excess return as the asset's own expected excess return, while the asset's contribution to the portfolio variance is twice the covariance between the asset's return and the portfolio return. Thus, the general efficiency condition can be written as follows:

$$\frac{E(\tilde{r}_j - r)}{cov(\tilde{r}_j, \tilde{r}_p)} = \lambda, \text{ for all risky assets } j\text{=1, ... N}, \tag{19.18}$$

where r is the risk-free rate of return, and $\tilde{r}_j$ the uncertain return on asset j. The common return/risk ratio, λ, depends on the investor's attitude toward risk, and is called the investor's *relative risk aversion.*

Example 19.4

Let there be two risky assets (j = 1, 2), with the following expected excess returns and covariances of return:

	$E(\tilde{r}_j - r)$	(co)variances	
Asset 1	0.092	$\text{cov}(\tilde{r}_1, \tilde{r}_1) = 0.04$	$\text{cov}(\tilde{r}_1, \tilde{r}_2) = 0.05$
Asset 2	0.148	$\text{cov}(\tilde{r}_2, \tilde{r}_1) = 0.05$	$\text{cov}(\tilde{r}_2, \tilde{r}_2) = 0.09$

Given these data, a portfolio p with weights ($x_1 = 0.4$, $x_2 = 0.6$) is efficient. We can verify the efficiency of this portfolio in two steps:

- First we compute the contribution of each asset to the total risk of portfolio p (covariance), as follows:[7]

  $$\begin{array}{ll} \text{Asset 1:} & \text{cov}(\tilde{r}_1, x_1\tilde{r}_1 + x_2\tilde{r}_2) = 0.4 \times 0.04 + 0.6 \times 0.05 = 0.046, \\ \text{Asset 2:} & \text{cov}(\tilde{r}_2, x_1\tilde{r}_1 + x_2\tilde{r}_2) = 0.4 \times 0.05 + 0.6 \times 0.09 = 0.074. \end{array}$$

- Next we compute, for each asset, the excess return/risk ratio and note that both ratios equal 2:

  $$\frac{0.092}{0.046} = 2 = \frac{0.148}{0.074}, \tag{19.19}$$

  which implies that the portfolio is efficient.

Moreover, this is not just any efficient portfolio: it actually is the tangency portfolio of risky assets. This is because (1) any efficient portfolio is a combination of the risk-free asset and the tangency portfolio of risky assets, and (2) this particular efficient portfolio contains no risk-free assets.

[7] We use the fact that the return on the risk-free asset does not co-vary with any risky asset's return.

The portfolio in the example will be selected by an investor with relative risk aversion equal to $\lambda = 2$. One way to detect differences in risk aversion among mean-variance investors is to watch the proportions they invest in the risk-free asset. An investor with a higher relative risk aversion simply allocates more of his or her wealth to the risk-free asset, and less to the tangency portfolio of risky assets.

Example 19.5

Suppose that an investor invests half of his or her wealth in the tangency portfolio identified in the previous example, and the remainder in the risk-free asset. That is, the weights in portfolio p' are $x_0 = 0.5$ for the risk-free asset, and ($x_1 = 0.2$, $x_2 = 0.3$) for the risky assets. We can easily verify that p' is still an efficient portfolio and that this investor has a relative risk aversion equal to 4:

- The risks of the assets in portfolio p' are computed as follows:

$$\begin{array}{ll} \text{Asset 1:} & \mathrm{cov}(\tilde{r}_1,\, x_1\tilde{r}_1 + x_2\tilde{r}_2) = 0.2 \times 0.04 + 0.3 \times 0.05 = 0.023, \\ \text{Asset 2:} & \mathrm{cov}(\tilde{r}_2,\, x_1\tilde{r}_1 + x_2\tilde{r}_2) = 0.2 \times 0.05 + 0.3 \times 0.09 = 0.037. \end{array}$$

- The excess return/risk ratios now both equal 4:

$$\frac{0.092}{0.023} = \frac{0.148}{0.037} = 4. \tag{19.20}$$

  which implies that the portfolio is also efficient.

Thus, the investor's relative risk aversion can be inferred from his or her portfolio choice. Relative to the tangency portfolio chosen by an investor with $\lambda = 2$, the more risk-averse investor with $\lambda = 4$ simply reduces the proportion invested in the risky assets by half. This, as we notice, also halves the (covariance) risks of each risky asset in the total portfolio. This stands to reason: if the total portfolio risk falls, assets' contributions to that total risk must fall too.

There is another, related, way to measure risk aversion: compute the excess-return-to-variance ratio for the entire portfolio. This ratio produces the same number as the previous ones since it takes the same linear combination of both numerators and denominators:[8]

$$\text{if } \frac{0.092}{0.023} = \frac{0.148}{0.037} = 4 \text{ then } \frac{0.2 \times 0.092 + 0.3 \times 0.148}{0.2 \times 0.023 + 0.3 \times 0.037} = 4. \tag{19.21}$$

We conclude that, for efficient portfolios, the holder's relative risk aversion can be measured by the overall excess-return/risk ratio:

$$\text{Relative risk aversion } = \lambda = \frac{\mathrm{E}(\tilde{r}_p - r)}{\mathrm{var}(\tilde{r}_p),}, \tag{19.22}$$

[8] The general way to establish this is to write the efficiency condition as $\mathrm{E}(\tilde{r}_j - r) = \lambda\, cov(\tilde{r}_j,\, \tilde{r}_p)$. This implies $x_j\mathrm{E}(\tilde{r}_j - r) = \lambda\, x_j\mathrm{cov}(\tilde{r}_j,\, \tilde{r}_p)$ and therefore $\sum_j x_j\mathrm{E}(\tilde{r}_j - r) = \lambda \sum_j x_j\mathrm{cov}(\tilde{r}_j,\, \tilde{r}_p) = \lambda\,\mathrm{cov}(\sum_j x_j\tilde{r}_j,\, \tilde{r}_p)$. Thus, $\mathrm{E}(\tilde{r}_p - r) = \lambda\, cov(\tilde{r}_p,\, \tilde{r}_p) = \lambda\,\mathrm{var}(\tilde{r}_p)$.

a relation that comes in good stead to derive the CAPM in the next subsection.

Using a variety of proxies for the market portfolio and a variety of methodologies, [19.22] has been used to estimate the US average risk aversion. The estimates vary, but the consensus in long-term tests is that λ exceeds unity. Also this result will come in handy later.

19.2.5 The Market Portfolio as the Benchmark

Let us now go from an individual investor's portfolio to the market portfolio—defined as the aggregate asset holdings of all investors in a particular group. The group typically considered in the standard CAPM is composed of all investors in the economy. What exactly "the" economy corresponds to in practice—a country? a region?—is left vague, but, crucially, this set of investors is assumed to have *homogeneous opportunities,* that is, equal access to the same list of assets, and *homogeneous expectations*, that is, equal perceptions about the return characteristics of the assets.

The effect of these homogeneity assumptions is that all of the investors agree about the composition of the tangency portfolio. If each investor holds the risk-free asset plus the same tangency portfolio, then also the aggregate portfolio must be a combination of the risk-free asset plus that very same tangency portfolio. But any such combination is efficient. Therefore, for the market portfolio (denoted by subscript m), the efficiency condition Equation [19.1] must hold, with λ_m now defined as the market's risk aversion (which can be shown to be a kind of weighted average of the individuals' risk aversions):

$$\frac{E(\tilde{r}_j - r)}{\mathrm{cov}(\tilde{r}_j, \tilde{r}_m)} = \lambda_m, \text{ for all risky assets } j\text{=1, ... N.} \tag{19.23}$$

Although Equation [19.23] is not yet written in the standard CAPM form, this equation already is an embryonic capital asset pricing model because it tells us what the expected excess return should be as a function of the asset's covariance risk in the market portfolio. To implement the model, we need to know the relative risk aversion for the average investor. But we just found a way to infer this: just use [19.22] to identify the market's relative risk aversion. This leads us straight to the CAPM:

$$\begin{aligned} \mathrm{E}(\tilde{r}_j - r) = \lambda_m \,\mathrm{cov}(\tilde{r}_j, \tilde{r}_m) &= \frac{\mathrm{E}(\tilde{r}_m - r)}{\mathrm{var}(\tilde{r}_m)} \mathrm{cov}(\tilde{r}_j, \tilde{r}_m) \\ &= \beta_{j,m} \,\mathrm{E}(\tilde{r}_m - r), \end{aligned} \tag{19.24}$$

In Equation [19.24], $\beta_{j,m} = \mathrm{cov}(\tilde{r}_j,\, \tilde{r}_m)/\mathrm{var}(\tilde{r}_m)$ is the asset's rescaled covariance risk, or the asset's beta. The advantage of rescaling the covariance risk is that $\beta_{j,m}$ is also the slope coefficient from the so-called *market model*, the regression of the return from asset j, on the return from the market portfolio, $\tilde{r}_j = \alpha_{j,m} + \beta_{j,m}\tilde{r}_m + \epsilon_{j,m}$.

Thus, the rescaled risk (the asset's relative risk, or market sensitivity) in Equation [19.24] can be estimated using time-series data of past stock returns and market returns, assuming, at least, that beta risks and expected returns are constant. We can summarize this model as follows:

- The beta is a measure of the asset's relative risk—that is, the asset's market covariance risk $\text{cov}(\tilde{r}_j, \tilde{r}_m)$, rescaled by the portfolio's total risk, $\text{var}(\tilde{r}_m)$. Beta can be estimated from the market-model regression.

- A risky asset with beta equal to zero should have an expected return that is equal to the risk-free rate, even if the asset's return is uncertain. The reason is that the marginal contribution to the total market risk is zero.

- If an asset's beta or relative risk is non-zero, the asset's expected return should contain a risk premium. The additional return that can be expected per unit of beta is the market's expected excess return above the risk-free rate.

19.2.6 A Replication Interpretation of the CAPM

An enlightening joint interpretation of the market model regression and the CAPM is as follows. A regression $\tilde{y} = a + b\tilde{x} + \tilde{e}$ has the property that it offers the best possible fit between $\tilde{y}$ and $a + b\tilde{x}$, in the sense that no other numbers a and b produce a lower residual variance, $\text{var}(\tilde{e})$. Now suppose that you were asked to find a combination of investments in the risk-free asset and a market-index fund that best resembles a particular asset, say Apple Computer common stock. This best-replication portfolio can be identified by regressing Apple's return onto the market return:

Example 19.6

Suppose that $\beta_{Apple} = 0.75$. If we invest 75 percent in the market and 25 percent in the risk-free asset, we hold a portfolio that offers the best possible replication of Apple Computer's return, among all portfolios that consist only of the market portfolio and the risk-free asset.

As, in the best replication, a fraction β is invested in the market and $(1 - \beta)$ in the risk-free asset, the expected return on such a best-replication portfolio would be

$$\begin{aligned} \text{E}(\tilde{r}_{Apple's\,replication}) &= \beta_{Apple}\,\text{E}(\tilde{r}_m) + (1 - \beta_{Apple})r \\ &= r + \beta_{Apple}\text{E}(\tilde{r}_m - r). \end{aligned} \tag{19.25}$$

But this is exactly the CAPM's prediction of the return on Apple itself. So the CAPM tells us that the expected return on stock j is equal to the expected return on its best replicating portfolio.

In that sense the logic of the CAPM is to some extent similar to the logic of asset-pricing-by-replication, as used in Part II of this book, except that we now use the best possible replication rather than exact replication. Because the replication is not exact, we need the CAPM assumptions to justify why the expected return on an asset should still be the same as the expected return on its best replicating portfolio, and why the market portfolio is the only replication instrument that is to be considered. In the CAPM logic, investors do not care about the imperfections in the replication (that is, the part of Apple's return not "explained" by the market) because they all hold the market portfolio anyway; the part of Apple's return not correlated with $\tilde{r}_m$ is simply diversified away.

19.2.7 When to Use the Single-Country CAPM

The CAPM as derived in Section 19.2 is routinely used in capital budgeting to determine the return that shareholders expect on investments with a given level of beta risk. For many countries, financial institutions provide estimates of the betas for various industries. Yet, one ought to interpret these figures with some caution. The assumption that underlies many of these estimates is that the CAPM holds country-by-country, in the sense that the market portfolio is equated with the portfolio of all assets issued by firms from that country alone. For example, beta service companies in the US tend to compute the beta of, say, the US computer industry by regressing the returns from a portfolio of US computer firms on the Vanguard index, which is an index of thousands of US stocks traded on the New York Stock Exchange, Amex, and NASDAQ. Likewise, in France, one would often estimate the risk of, say, the French steel industry by regressing the returns from a portfolio of steel companies on the index of French stocks. In the same vein, the expected excess return on the market would be estimated from past returns on the Vanguard index or on the index of French stocks traded at the Paris section of Euronext, respectively.

Is the market portfolio of assets held by a country's investors the same as the portfolio of assets issued by the country's corporations? This is only true if investors have access to local shares only *and* all local shares are held by residents of the country. That is, if one equates the market portfolio with the portfolio of locally issued shares, capital markets are assumed to be fully segmented. However, in most countries there are no rules against international share ownership; investors can easily diversify into foreign assets, and foreigners are allowed to buy domestic shares. Thus, the traditional interpretation that the market portfolio consists of the index of stocks issued by local companies is valid only in segmented markets.

Example 19.7

Until the later 1990s, the stock markets of India, South Korea, and Taiwan were almost perfectly segmented from the rest of the world in the sense that foreigners could buy only a small fraction of the local stocks, and local investors could not easily buy foreign assets. Thus, the Indian market portfolio was essentially the same as the portfolio of stocks issued by Indian firms, and similar for Korea and Taiwan.

In the presence of market segmentation, the cost of capital to be used by, say, a

North-American or European firm is likely to be different from the cost of capital to be used by an Indian firm, even when these companies are evaluating similar investments. For the Indian case, we would have used a one-country CAPM. The question addressed in the next section is how, say, a Canadian firm should determine its cost of capital, knowing that its investors are part of a market that is much wider than just Canada. There are no rules that prevent Canadian investors from buying US or European assets, nor are nonresidents barred from buying Canadian stocks. Under these circumstances the index of stocks issued by Canadian firms is likely to be a poor proxy for the portfolio held by the average Canadian investor. It follows that a Canadian firm cannot use the single-country CAPM to set the cost of capital for an investment project. Not only does the Canadian-stock index miss foreign stocks held by residents, but it also ignores the fact that many Canadian stocks are held by foreigners. Note also that this problem arises whether the project is domestic or foreign: it's not as if Canadians can still use a one-country CAPM for home investments, and only have a problem if the project is foreign.

19.3 The International CAPM

As we just stated, there are no rules preventing Canadian investors from buying US or European assets; nor are there any regulations barring nonresidents from buying Canadian stocks. Still, this mere fact is not sufficient to lead to international diversification by investors. We have already argued, in Chapter 18, that there are strong incentives for investors to diversify internationally. We just pointed out why this causes a problem with the standard CAPM, at least in the version that uses the locally-issued stock index rather than the locally-held stock index. From this starting point we add four items: we explain the role of exchange risk for asset pricing in an internationally integrated capital market; we derive a two-country version of the International CAPM of Solnik (1973) and Sercu (1980, 1981); we generalize to the case with many countries and stochastic inflation; and we conclude with a review of empirical tests of the International CAPM.

19.3.1 International diversification and the traditional CAPM

International diversification is beneficial for the investor, and investors do use this added opportunity to reduce risks. Clearly, it is then no longer acceptable to use a CAPM equation with, as its benchmark portfolio, the local stock index (defined as the index of all securities issued by firms incorporated in the country). First, this benchmark omits foreign assets, which represent an important component of the local investor's asset holdings. Second, this benchmark ignores the fact that a substantial part of the stocks issued by local corporations are, in fact, held by nonresidents.[9] All of this means that, in internationally integrated markets, the true stock market portfolio for any country is unobservable—and, with an unobservable national stock market portfolio, the standard CAPM is of no practical use to managers who, for instance, want to assess the cost of capital or evaluate the performance of their investment advisers.

19.3.2 Why Exchange Risk Pops up in the International Asset Pricing Model

How can we get around this problem of an unobservable market portfolio? One could argue that, even if we do not know what shares are held by whom, we can still observe the *world* market portfolio. (For conciseness, we will refer to the the countries that allow free capital movements as "the world", with an apology to residents from China and other unworldly countries.) Even if we do not know what stock is held by whom and where, we do know what stocks are listed somewhere in the world and how many shares are outstanding at what price. Thus, the world market portfolio contains all securities issued by all firms in the world, and it can be obtained by constructing a value-weighted sum of all member countries' local stock indices.[10] As investors do hold assets from all over the world, and as the world market portfolio is observable, a very simple approach to international asset pricing would be to interpret the world as one huge country, and use the world market portfolio as the benchmark in a unified-world CAPM.

There is, however, one important reason why international asset pricing in integrated capital markets cannot simply be reduced to an as-if-one-country CAPM. Even if international capital transactions are unrestricted and have low costs, transactions in the commodity markets are still difficult and costly. These imperfections in the goods market, as we saw in Chapter 3, lead to substantial deviations from relative purchasing power parity and to real exchange risk. The (real) return on, say, IBM common stock as realized by a German investor differs from the (real) return realized by a Japanese investor on the same asset. As a result, the distributions of the real return from a given asset depend on the nationality of the investor. This then violates the homogeneous expectations assumption of the CAPM, which states that all investors agree on the probability distribution of the (real) asset returns. In a sense, the investors' perceptions about real return distributions are segmented along country lines because goods prices differ across countries, implying that investors

[9]The same problem arises when one includes into the market portfolio all stocks—domestic or foreign—that are listed on the national stock exchange(s). Investors can (and do) buy foreign assets in foreign stock exchanges, or can (and do) buy foreign assets through mutual funds that are traded over-the-counter; and all of these investments are missing from the menu of locally listed stocks.

[10]A well-known proxy for such an international stock market index is the Morgan Stanley Capital International (MSCI) index, or Datastream's World Market Index. Both are biased towards large firms; but small firms are held locally, mostly, so that's not a huge problem.

from various countries have different views on the distributions of real returns on any given asset or portfolio.

Example 19.8

A clear example is the return on the two countries' T-bills. Suppose that there is no inflation. While to a US investor, the CAD T-bill is one of the available risky assets, it is risk-free to a Canadian investor. On the other hand, the USD T-bill is a risky asset to a Canadian investor but risk-free to a US investor. Thus, the perceived distribution of (real) returns depends on the nationality of the investor.

Thus, we need to derive a CAPM that takes into account the heterogeneous viewpoints of investors from various countries. In keeping with our discussion of the standard CAPM, we initially ignore inflation. To simplify the analysis, we shall initially assume there are just two countries, the US and Canada. Once you understand the two-country model, you can easily generalize to the case of many countries.

The problem is that the Canadian investor's portfolio choice is based on how each asset contributes to the variance and expected excess return on the portfolio measured in CAD, while the US investor's portfolio choice is based on the assets' contributions to a portfolio whose risk and return are expressed in USD. Let, as usual, the asterisk refers to the foreign country (say, the US); p^* refers to the portfolio held by the US investor; $\tilde{r}_j^*$ refers to a return in FC on stock j (whose nationality, if any, we do not really need to know); r^*, unsubscripted, as usual refers to the USD risk-free rate; and $\tilde{r}_{p^*}^*$ denotes the return, in USD, on the US market portfolio p^*. Then what we know about portfolio choice can be summarized as follows:[11]

$$\text{Canadians choose } p \text{ such that } \mathrm{E}(\tilde{r}_j - r) = \lambda \operatorname{cov}(\tilde{r}_j,\, \tilde{r}_p), \tag{19.26}$$

$$\text{Americans choose } p^* \text{ such that } \mathrm{E}(\tilde{r}_j^* - r^*) = \lambda \operatorname{cov}(\tilde{r}_j^*,\, \tilde{r}_{p^*}^*). \tag{19.27}$$

What, then, is the relation between expected excess returns and the world market portfolio, which is the sum of p and p^*? To identify that link, we have to translate [19.27] into the same currency as [19.26], the CAD. Using a trick called Ito's Lemma (see Technical Note 19.2), [19.27] can be translated into CAD as follows:

$$\text{Americans choose } p^* \text{ such that } \mathrm{E}(\tilde{r}_j - r) = \lambda \operatorname{cov}(\tilde{r}_j,\, \tilde{r}_{p^*}) + (1-\lambda) \operatorname{cov}(\tilde{r}_j, \tilde{s}), \tag{19.28}$$

where $\tilde{s}$ is the percentage change in the exchange rate (CAD per USD). What is going on here is that US investors really care about their wealth expressed in USD, W_{us}^*, because the consumption prices relevant to them are (almost) constant in USD and far less so in CAD. We can always re-express W_{us}^* as CAD-measured wealth divided by the CAD/USD exchange rate, $W_{us}^* = W_{us}/S$. So people who care about W^* will act as if they care about wealth in CAD sure enough—because, everything else

[11] It would not have been very painful to allow for different risk aversions across countries too, but little additional insight would have been gained, so we set $\lambda^* = \lambda$.

Table 19.2: **Exchange Rate Exposure: Good or Bad?**

	Example 1: covariance > 0		Example 2: covariance < 0		Comment
W_{us} (in CAD)	12,000	16,000	12,000	16,000	same distribution for W_{us} ...
S (CAD/USD)	1.00	1.50	1.50	1.00	and same distribution for S;
W^*_{us} (in USD)	12,000	10,667	8,000	16,000	but the positive-cov case has ...
$\mathrm{E}(W^*_{us})$	11,333		12,000		– a lower mean W^*_{us} ...
stdev(W^*_{us})	667		4,000		– and a lower stdev W^*_{us}

being the same, the higher their CAD wealth, the higher also their wealth in USD. The fact that, holding constant the exchange rate, they care about CAD-expressed wealth then explains why the first half of the efficiency condition looks like the Canadian investor's condition. But US investors will all the time also think of the exchange rate, because deep down they care about USD-measured wealth only. It is this concern about the exchange rate that induces a second item. But, as we shall see, it is less obvious whether the US investor, thinking in CAD terms but caring about USD numbers, likes exchange-rate exposure or not.

Example 19.9

In Table 19.2 we have picked two examples where, in each example, there are two equally probable scenarios for CAD wealth and the exchange rate. The means and variances are the same across the two examples, but the first one has a positive association between CAD wealth and the exchange rate while in the second example the correlation is negative. We see that a larger positive covariance is a mixed blessing: it lowers both the mean (bad!) and the variance (good!). So whether on balance the effect is preferred depends on your degree of risk aversion, notably whether you attach more weight to the rise in return than to the rise of risk.

It can, in fact, be shown that investors with risk aversion equal to 1 ignore covariance with S. More risk-averse investors ($\lambda > 1$) like it because they like the variance-reduction effect, while less risk-averse people dislike it: the drop in the mean is viewed as too high a price for the lower risk. But note that, among financial economists, the consensus probably is that lambda exceeds unity. (Macro-economists are not so sure.) Thus, the modal investor probably prefers the hedging effect and is willing to accept a lower mean return on asset j if it does help as a hedge.

What assets would be especially attractive to US investors from that perspective? One might guess that US stocks may be more appealing than Canadian stocks. But such a view may be simplistic, as the next subsection argues.

19.3.3 Do Assets have a Clear Nationality?

For a better understanding of the exchange rate covariance risk of individual assets, it is convenient to scale the covariance risk by the exchange rate variance. Consider the following regression equation:

$$\tilde{r}_j = \alpha_{j,s} + \gamma_j\, \tilde{s}_{\text{CAD/USD}} + \epsilon_{j,s}. \tag{19.29}$$

The regression coefficient γ_j equals $\text{cov}(\tilde{r}_j, \tilde{s})/\text{var}(\tilde{s})$—the asset's exchange rate covariance risk, scaled by the variance of the exchange rate change. In this sense, γ_j measures the relative exchange risk of asset j, or the relative exposure of asset j to the exchange rate, in the same way beta measures the relative exposure of a stock to market movements. We now consider the exchange rate exposure of six types of assets: a domestic and foreign risk-free asset, a foreign exporter and importer, and a domestic exporter and importer.

- Let us consider the domestic T-bill, asset 0. Since this return is not stochastic, it has zero exposure to the exchange rate.
- The next asset we consider is the USD T-bill, asset 1. The return, measured in CAD, on the USD T-bill increases by one percent if the CAD/USD spot rate increases by one percent. This follows from

$$\tilde{r}_{\text{USD}\,Tbill} \approx r^* + \tilde{s}_{\text{CAD/USD}}. \tag{19.30}$$

  Clearly, if $\tilde{r}_j = r^* + \tilde{s}_{\text{CAD/USD}}$, then, in the relative exposure regression Equation [19.29], we must have $\gamma_{\text{USD}\,Tbill} = 1$ (and $\alpha_{\text{USD}\,Tbill} = r^*$). In this sense, the exposure regression (Equation [19.29]) for the foreign T-bill will reveal a very clear nationality for that asset. In CAD terms, the USD T-bill is exposed one-to-one to its "own" exchange rate, CAD/USD.

  Thus far, things are clear: the home T-bill has zero exposure and the foreign one has a unit exposure. For stocks, however, nationality is much more blurred:

- Let asset 2 be a Canadian importer. Typically an appreciation of the USD relative to the CAD is bad news for such Canadian firm, because its costs have gone up. Thus, for a Canadian importer, the relative exposure (γ_j) is negative.
- Let us now consider as asset 3 a Canadian producer competing against US producers in the US and/or Canadian market. Typically an appreciation of the USD relative to the CAD is good news for such a Canadian firm, because its competitive position has improved. Thus, for a Canadian exporter or import-substituter, the relative exposure (γ_j) is positive.
- The next case we look at is a US corporation that competes against Canadian firms in the US and/or Canadian market. Holding constant the USD price of the stock, a one percent appreciation of the USD adds one percent to the return on that stock in CAD. However, an appreciation of the USD simultaneously is

Figure 19.4: **Relative exposures (γ) of various assets**

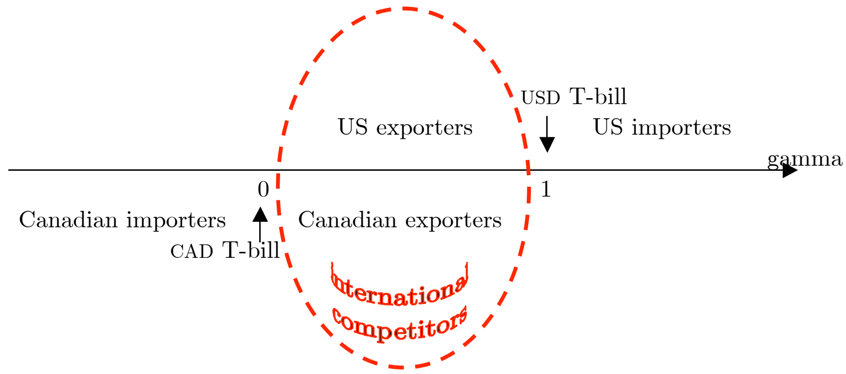

bad news for this company, because its competitive position has deteriorated. Thus, the price of the stock measured in USD typically drops when the USD appreciates. This drop in the USD value of the stock weakens the effect of the exchange rate itself, and will lead to a relative exposure that is below unity.

Example 19.10

Suppose that, empirically, the stock price in USD of a US firm goes down by, on average, 0.25 percent for a 1 percent increase in the CAD/USD rate. This then implies that the return, in CAD, on the stock will go up by about 0.75 percent for a one percent rise in the CAD/USD rate. That is, the Canadian investor on average suffers a 0.25 percent capital loss in USD terms, which is to be subtracted from the 1 percent gain on the USD itself.

- Lastly, consider a US importer. An appreciation of the USD relative to the CAD is good news for this US firm, because its costs have gone down. Thus, for a US importer, we would typically see a rise of the USD stock price, reinforcing the effect that the exchange rate itself has on the asset's CAD value. Thus, the gamma would exceed unity.

We conclude that exchange rate covariance risks can be very different for different assets. The relative exposure of a foreign T-bill is unity, but the relative exposure of a foreign stock could be higher, or lower. Notably, there is a whole group of foreign firms with gamma's below 1, and a bunch of domestic firms with gamma's above 0. We'd probably better speak of all of these as internationally competing firms that do not fundamentally differ from each other. In short, unlike T-bills, their stocks have no clear-cut economic nationality.

19.3.4 The International CAPM

Let us again consider the two equations that determine the Canadian and US market portfolios:

$$\text{CDN: } \mathrm{E}(\tilde{r}_j - r) = \lambda \operatorname{cov}(\tilde{r}_j, \tilde{r}_p), \tag{19.31}$$

$$\text{US: } \mathrm{E}(\tilde{r}_j - r) = \lambda \operatorname{cov}(\tilde{r}_j, \tilde{r}_{p^*}) + (1-\lambda) \operatorname{cov}(\tilde{r}_j, \tilde{s}). \tag{19.32}$$

In Technical Note 19.3 it is shown that these equations can be aggregated into the following:

$$\mathrm{E}(\tilde{r}_j - r) = \lambda \operatorname{cov}(\tilde{r}_j, \tilde{r}_w) + \kappa \operatorname{cov}(\tilde{r}_j, \tilde{s}), \tag{19.33}$$

with $\tilde{r}_w$ referrring to the return on the world market portfolio and κ being a function of the national invested wealths and the national (unity minus) risk aversions. Compared to the country-by-country efficiency conditions, what we now have on the right-hand side is a covariance with the *world* market portfolio, which is more observable than the national portfolios, and a covariance with exchange rate, the result of taking into account the heterogeneous expectations induced by exchange rate uncertainty.

This is, again, half a CAPM in the sense that it tells us what expected returns should be, taking into account the risks of the assets. As before, we need to know the prices of risk before this is of any use whatsoever to an investor or analyst. The approach is the same as before except that we now need two benchmarks. If we pick the world market portfolio and the USD treasury bill, a simple generalization of the one-country CAPM emerges, as shown in Technical Note 19.4:

$$\mathrm{E}(\tilde{r}_j - r) = \beta_{j,w}\mathrm{E}(\tilde{r}_w - r) + \gamma_{j,s}\mathrm{E}(\tilde{s} + r^* - r), \tag{19.34}$$

where beta and gamma are from the multiple regression that combines the market model and the exposure model we considered in the preceding subsection:

$$\tilde{r}_j = \alpha_{j,w,s} + \beta_{j,w;s}\tilde{r}_w + \gamma_{j,s;w}\tilde{s} + \tilde{\epsilon}_{j;w,s}. \tag{19.35}$$

The subscript $j;s$ to beta intends to remind you that this is not the simple beta we are used to: we are now holding constant the exchange rate. Likewise, the subscript $j;w$ to gamma tells you we are now holding constant the world market return, unlike in the simple exposure regression we looked at a few pages up.

To interpret the regression [19.35] and the International CAPM [19.34], note that the regression again identifies the best possible replication of asset j that one can achieve using the two benchmark portfolios, the world market portfolio and the foreign T-bill, along with the risk-free asset.

Example 19.11

Suppose that, for a US stock, the coefficients in Equation [19.35] are estimated as $\beta_{j,w;s} = 1.2$ and $\gamma_{j,s;w} = 0.75$. Consider portfolios that consist of an investment in the world market portfolio (with weight x_w), an investment in the USD T-bill (with weight x_s), and weight $1 - x_w - x_s$ invested in the CAD risk-free asset. If $\beta_j = 1.2$ and $\gamma_j = 0.75$, we invest $x_w = 1.2$ in the world market portfolio, $x_s = 0.75$ in the USD T-bill, and 1 – 1.20 – 0.75 = –0.95 in the domestic risk-free asset. This portfolio provides the best possible replication of the return from asset j using just the two benchmark portfolios as replicating instruments.

The International CAPM then says that the expected return on a stock j is the same as the expected return on the stock's best replication portfolio—see Technical Note 19.5 for the details:

Example 19.12

Continue the same example ($\beta_{j,w;s} = 1.2$ and $\gamma_{j,s;w} = 0.75$). If the world market portfolio has an estimated risk premium of 0.05 and the currency of 0.01 *p.a.*, then the expected risk premium on the stock is estimated as $1.2 \times 0.05 + 0.75 \times 0.01 = 0.0675$, or 6.75 percent (on top of the risk-free rate).

19.3.5 The N-Country CAPM

The "world" (in the sense of the integrated capital market) has far more countries than two. The generalisation of the two-country model is obvious. First, there will be as many gamma's as there are exchange rates in the world. Second, the beta and the gammas must be estimated from one regression containing r_w and all the $\tilde{s}_i$'s:

$$\mathrm{E}(\tilde{r}_j - r) = \beta_{j,w;..}\mathrm{E}(\tilde{r}_w - r) + \gamma_{j,s_1;..}\mathrm{E}(\tilde{s_1} + r_1^* - r) + \gamma_{j,s_2;..}\mathrm{E}(\tilde{s_2} + r_2^* - r) + ...\gamma_{j,s_n;..}\mathrm{E}(\tilde{s_n} + r_n^* - r), \tag{19.36}$$

where beta and the n gammas are from the multiple regression that combines the market model and n exposure models, one per currency, that we considered in the preceding subsection:[12]

$$r_j = \alpha_{j,w;\underline{s}} + \beta_{j,w;..}r_w + \gamma_{j,s_1;..}\tilde{s_1} + \gamma_{j,s_2;..}\tilde{s_2} + ...\gamma_{j,s_n;..}\tilde{s_n} + \epsilon_{j;w,\underline{s}}. \tag{19.37}$$

In practical applications, restraint is recommendable, as Goethe would readily concur. A CAPM *cum* regression of 150 terms will not do: it would add more noise than information. One reason is that exchange-risk premia $\mathrm{E}(\tilde{s} + r^* - r)$ are empirically small, have a long-run mean that is hard to statistically distinguish from zero, and are not easy to estimate with reasonable precision. Also, gammas are similarly difficult to estimate precisely. So my advice is to surely restrict, *a priori*, the list of countries to those where there is a good common-sense reason for expecting an exposure, and censor away the gammas with the wrong size or

[12] Apologies for the baroque subscripts. The semi-colon usually initiates a list of variables that are held constant. Here the list would be too long, so we drop it. Still, you should remember that these are multiple-regression coefficients, measuring the impact of one variable holding constant the other ones.

sign. Personally I would perhaps even entirely omit the exposure terms: given the uncertainties surrounding the risk premia and the exposures, one might just work with the world-market term in the i-CAPM, and simply widen the scope of the sensitivity analysis that should be part and parcel of every capital-budgeting exercise:

$$\mathrm{E}(\tilde{r}_j - r) \approx \beta_{j,w;..}\mathrm{E}(\tilde{r}_w - r). \tag{19.38}$$

The only surviving difference with the standard CAPM would then be the use of a world market as benchmark, and the multiple beta.[13]

19.3.6 Empirical Tests of the International CAPM

In this chapter, we are suggesting that you replace your familiar single-market CAPM equation by a more complicated version, Equation [19.36] or [19.38]. The first issue is whether one of the basic assumptions of the international model, the absence of controls on capital flows, is reasonable. Second, are the empirical data compatible with the International CAPM and, if so, can we also reject the single-country view of the world?

Let us first examine the effect of direct controls on foreign investment. The controls may either limit foreign investment into a country or restrict domestic residents from investing abroad. Restrictions on foreign investment into a country may be imposed in different ways—in the form of a limit on the fraction of equity that can be held by foreigners or a restriction on the types of industries in which foreigners can invest. Historical details on the type and magnitude of these restrictions can be found in Eun and Janakiramanan (1986, Table 1). There may also be domestic controls on how much a resident can invest abroad. For example, Japanese insurance companies could not invest more than 30 percent of their portfolio in foreign assets at the time, and only 30 percent of Spanish pension funds could be invested abroad. Two questions need to be answered. One, if these restrictions exist, do they have a significant impact on the choice of the optimal portfolio and hence, potentially, on asset pricing? Two, how significant are these constraints today?

Bonser-Neal, Brauer, Neal and Wheatley (1990) examine whether the restrictions on investing abroad are binding. They look at closed-end country funds and find that these mutual funds trade at premia relative to their net asset values, indicating

[13] The need to still use a multivariate regression even in the truncated model follows from the fact that our basic model is Equation [19.33], not Equation [19.23]. Equation [19.33] simplifies to the univariate equation, [19.23], only if either the prices of exchange covariance risk, η_k, are all zero, or the covariances between asset returns and exchange rate changes themselves are zero. The first case requires very special utility functions (with $\lambda = 1$), and the second case cannot possibly be true for all assets and home currencies simultaneously. Thus, we do need the multivariate model. Moreover, although the risk premium for exchange risk can be small it is unlikely to be exactly zero. That is, we use the one-factor world model merely as an approximation. If we would, in addition, use a univariate beta, we would introduce another (unnecessary) error to the approximation.

that the French, Japanese, Korean, and Mexican markets are at least partially segmented from the US capital market. Hietala (1989) studies the effects of the Finnish law that prevented investors from investing in foreign securities and finds that there is a significant difference between the returns on domestic assets required by residents compared to foreigners. Gultekin, Gultekin, and Penati (1989) find strong evidence that the US and Japanese markets were segmented prior to 1980. However, while there were substantial controls on capital flows before the 1980s, this is no longer true. Halliday (1989) already reports that even in those days there were few constraints on investing in foreign stock markets. This was and is especially true for investing in the markets of developed countries. For example, already in the 1980s there were no controls on international investment into or from Austria, Belgium, Denmark, Ireland, Italy, Japan, the Netherlands, the UK, the US, and West Germany. The controls studied by Hietala (1989) and Gultekin, Gultekin, and Penati (1989) were removed in 1986 and 1980, respectively. Also, looking at restrictions that limit domestic residents from investing abroad, one sees that these constraints are often not binding. For example, Fairlamb (1989) reports that in 1988 only 8 percent of Spanish funds were actually invested in foreign assets, while the limit was 30 percent. Thus, while direct controls on foreign investment may have been important in the past, they are probably no longer an important determinant of portfolio choice and asset pricing in the main OECD countries.

Let us now discuss the more direct tests of international asset pricing models. Solnik (1973), who did the first theoretical and empirical work in international asset pricing, tests a special case of Equation [19.36], where the world market risk premium and exposure risk premia could be merged into one single term. He concludes that the data are consistent with his International CAPM, although he does not test his model against the single-country alternative.

An early test that does compare an international asset pricing model against the single-country alternative was carried out by Stehle (1976). Specifically, Stehle tries to find out empirically whether US stocks are priced in a national market or in a world market. He, too, uses a restricted version of Equation [19.36], assuming that λ equals unity so that all currency risk premia disappear. The only remaining difference between the international model (Equation [19.36]) and the national model is the definition of the market portfolio. Specifically, in Equation [19.36], the benchmark portfolio is the world market portfolio, while in Equation [19.24], it is the national market portfolio. Stehle's tests are not able to empirically reject one in favor of the other, and Stehle concludes that asset pricing is done in a single-market context. Dumas (1976), however, argues that when the data do not allow one to distinguish between single-country asset pricing and international asset pricing, then one ought to retain the simplest view—that is, one should conclude that there is one international market instead of the many separate national markets.

There have been many additional empirical investigations, with a large portion of them testing special restricted versions of Equation [19.36]. The conclusions tended to be ambiguous. But more recent work has come up with more definite answers. As already mentioned, Gultekin, Gultekin, and Penati (1989) provide strong evidence that the US and Japanese markets were segmented prior to 1980. However, they also show that after the enactment of the Foreign Exchange and Foreign Trade Control law in 1980, there is no longer any significant evidence against the hypothesis that US and Japanese stocks are priced in an integrated market. A careful, and more recent, test is by Dumas and Solnik (1991). They test the Solnik-Sercu International CAPM, allowing for changes in risks and risk premia over time. Using data from major OECD countries, they reject Stehle's hypothesis that the exposure risk premia, γ_i, are zero, but they do not reject the full version (with non-zero risk premia for exchange rate exposure). They also reject single-country asset pricing (with a purely local benchmark). All of this lends support to the International CAPM, at least for the major OECD countries that do not impose explicit restrictions on capital movement. There are also a few papers by De Santis and Gerard (1997, 1998) that allow for autocorrelation in not just expected returns but also in variances and covariances, modeling the fact that risk comes in waves. Their work confirms that exchange-rate exposure is often non-zero and earns a statistically significant premium.

19.4 The CFO's Summary *re* Capital Budgeting

International asset pricing is potentially complicated by two extra issues: exchange risk, and segmentation of capital markets. If the capital market of the home country and the host country are integrated, the cash flows of an investment project can be valued in any currency, including the host currency. This simplifies capital budgeting in the sense that no exchange rate forecasts seem to be necessary for the translation. On the other hand, in integrated markets it becomes impossible to observe the portfolio of risky assets held by the average investor in any of the individual countries. Thus, the International CAPM has to be used, which means that, in principle, exchange rate expectations and exposures still show up in the cost of capital. In short, forecasts and exposures can only be eliminated by cutting corners.

Thus, the first issue is whether or not there is integration. Having selected either the single-country CAPM or the International CAPM, the next issue is to obtain estimates of the model parameters. We need the stock market sensitivity or beta and, in the International CAPM, the exchange rate exposures. We also need the expected return on the corresponding benchmark portfolios.

19.4.1 Determining the Relevant Model

If the capital market of the home country and the host country are segmented from each other, the investing firm should set the cost of capital equal to the return that is expected by its own shareholders. This means that a particular investment may be profitable for a foreign firm but not profitable for a local firm.

Table 19.3: **Rules for the Capital-Budgeting Process: Overview**

	CoCa model	currency of calculations
1. Foreign investments:		
• home and host financially integrated	inCAPM	FC & HC
• home and host financially segmented		
• home country part of larger financial market	inCAPM	HC only
• home country totally isolated	CAPM	HC only
2. domestic investments		
• home country part of larger financial market	inCAPM	n.a.
• home country totally isolated	CAPM	n.a.

Example 19.13

At the time of writing, the Chilean stock market remains strongly segmented from the rest of the world. If a Chilean firm makes an investment in Chile, the firm will estimate the beta by regressing returns from a portfolio of stocks in the same industry on the Chilean stock market index. Note that the returns from this investment are likely to be strongly correlated with the Chilean market index because there are important common factors, like the business cycle or interest rates, that affect all Chilean firms in similar ways. Thus, the investment is relatively risky for an Chilean firm. But the same project may be low-risk from the point of view of, say, an Austrian firm. The reason is that, because the Chilean economy is only loosely connected to North-America, Europe and Asia, the returns from the Chilean project will not be highly correlated with the returns on the typical world investor's portfolio, which is strongly diversified internationally. So the investment adds little to the risk of an international portfolio, but much more to the risk of a purely Chilean portfolio.

Note that segmentation of the home-country and the host-country capital markets does not mean that each market is a single-country market. The shareholders of the Austrian firm are likely to live in many different countries, and they all have access to non-Austrian shares, too. Thus, it is appropriate for the Austrian firm to set its cost of capital using an international model, that is, using the "world" market portfolio as a proxy for the true benchmark relevant to its shareholders.

19.4.2 Estimating the Risk of a Project

The market risk and the exchange risk exposures are defined as the slope coefficients in the regression of j's return on the world market return and all relevant exchange rate changes. Estimates obtained from time series of past data are subject to substantial estimation errors, stemming from pure sample-specific coincidences.

A standard solution is to estimate the risks from returns on industry portfolios rather than from individual stock data. That is, one estimates returns on, typically, an equally weighted portfolio of all stocks in the same industry i: One then estimates the risks by regressing industry-portfolio returns rather than individual stock returns. The underlying idea is that, as portfolio returns are more diversified, there is less residual noise in the regression, which improves the quality of the estimates.

Example 19.14

Suppose that Toyota considers building a new plant in the UK, which would sell its output in the entire European Union. Then Toyota could estimate the beta and gammas of the European car industry as a whole, rather than estimating the risks using just a simple stock.

Still, the portfolio approach assumes that all firms in the index have the same risks. In practice, one would often have serious difficulties in identifying a sufficiently large number of firms that have the same exposure as the project at hand.

Example 19.15

Suppose that Oerlikon, a Swiss firm, wants to build a plant for the production and sale of maintenance welding electrodes in India. There may be a number of Indian firms active in the welding industry, but not one of them is priced in the OECD capital market. Hence, Oerlikon cannot directly measure the risk of the Indian welding industry relative to the world market portfolio.[14] Thus, when valuing the project, Oerlikon would have to use an indirect, forward-looking approach to assess the risk. For instance, Oerlikon could argue that (1) the maintenance welding industry is not very cyclical, (2) the Indian business cycle is still largely independent of economic cycles in the OECD, so that (3) the beta of this Indian project relative to the OECD market portfolio is bound to be low. In addition, Oerlikon could argue that the exposures of Rupee cash flows to OECD exchange rates are small or zero because the Indian economy is still relatively closed. In short, beta is probably low; the Rupee gamma is probably equal to unity or thereabouts (as cashflows are unexposed in Rupee terms); and the other gammas must be close to zero.

Data availability is just one possible issue. The relevance of any available data is another. As pointed out in Chapter 13, exchange risk exposure when you are at the top of a PPP-deviation cycle would be very different from an exposure when the currency is at a low, in real terms. In such case, rather than estimating a misleading gamma you could (*i*) work with forward-looking scenarios, see Chapter 13 and then hedge the currency effect on the basis of the implied exposure; or (*ii*) ignore currency elements in the cost of capital, and widen the range of the sensitivity analyses.

[14] A procedure that consists of translating rupee returns on Indian stocks into an OECD currency and then estimating the risks is flawed because the prices of these Indian companies in the Bombay stock market are different from what they would have been if the assets had been priced internationally.

19.4.3 Estimating the Risk premia

Assuming that we have an approximate idea of the beta and gammas, we need estimates of the expected risk premia per unit of risk. The expected excess return on the world market portfolio is still rather hard to estimate, even though it is not quite as bad as a typical currency-risk premium. The sample averages of returns observed in the past differ substantially across sample periods, and it is also known that the expected return changes over time.[15] Still, we know that there is a positive risk premium on the world stock market, and variations over time in the expected excess return may not be overly important when the NPV evaluation horizon is, say, one decade rather than a month or two days.

Turning to the expected excess return on the various foreign T-bills, these risk premia also change over time, as we have seen in Chapter 10—and, unlike for the world market risk premium, we are not even sure whether the long-run mean actually differs from zero. Since exchange risk premia are small in the short run and close to zero in the long run, for practical applications one might have to be content with an approach that ignores these and use the following simplified version of Equation [19.36]:

$$\mathrm{E}(\tilde{r}_j - r) \approx \beta_{j,w;\underline{s}}\mathrm{E}(\tilde{r}_w - r), \tag{19.39}$$

where the beta is still estimated from a multivariate regression (Equation [19.37]) rather than from a bivariate regression).

You should not be overly discouraged by these approximations. No model is perfect; and the International CAPM does work better than competing models. Still, the cost of capital is measured imperfectly, and NPV computations should always be undertaken for a whole range of reasonable discount rates, to see to what extent the accept/reject recommendation is sensitive to the estimate of the cost of capital.

[15]The return is partially predictable on the basis of (1) the risk spread (the difference between low-grade bond yields and government bond yields), (2) the term spread (the difference between short-term and long-term bond yields), and (3) the dividend yield.

19.5 Technical Notes

Technical Note 19.1 The efficiency condition

Let the desirability of the portfolio p be denoted by $V_p = V(\mathrm{E}(\tilde{r}_p - r), \mathrm{var}(\tilde{r}_p))$. The optimum is found by setting, for each risky asset j, the derivative of V_p w.r.t. x_j equal to zero. The effect of a small change in x_j on V_p works through two channels: the expectation, and the variance; so below we see x_j's effect on V_p via the mean, and similarly x_j's effect on V_p via the variance. In the second line we fill in the effect of x_j on mean and variance, Equations [19.16] and [19.17]:

$$\begin{aligned} 0 &= \frac{\partial V}{\partial x_j} = \frac{\partial V}{\partial \mathrm{E}()}\frac{\partial \mathrm{E}()}{\partial x_j} + \frac{\partial V}{\partial \mathrm{var}()}\frac{\partial \mathrm{var}()}{\partial x_j}, \\ &= \frac{\partial V}{\partial \mathrm{E}()}\mathrm{E}(\tilde{r}_j - r) + \frac{\partial V}{\partial \mathrm{var}()}2\mathrm{cov}(\tilde{r}_j, \tilde{r}_p); \\ \Rightarrow 0 &= \mathrm{E}(\tilde{r}_j - r) - \lambda_p \mathrm{cov}(\tilde{r}_j, \tilde{r}_p), \end{aligned} \tag{19.40}$$

where $\lambda_p := -2\frac{\partial V/\partial \mathrm{var}()}{\partial V/\partial \mathrm{E}()}$. This is a positive number since a higher variance lowers the desirability V while a higher expected return increases it. Crypto-mathematicians recognize this ratio of partial derivatives as the implicit derivative (or marginal trade-off) of mean for variance in the chosen solution.

Technical Note 19.2 Using Ito's Lemma to transcribe the fc efficiency condition.

Start by relating the CAD return on j to the USD return: $1 + \tilde{r}_j = (1 + \tilde{r}_j^*)(1 + \tilde{s})$, with $\tilde{s} = \Delta S/S$ and S is CAD/USD. Solve for $\tilde{r}_j^*$ and Taylor-expand as follows:

$$\tilde{r}_j^* = \frac{1 + \tilde{r}_j}{1 + \tilde{s}} - 1 \approx \tilde{r}_j - \tilde{s} - [\tilde{r}_j\, \tilde{s}] + \tilde{s}^2, \tag{19.41}$$

A readily acceptable result of Ito's Lemma is that, for shorter and shorter holding periods, products of three or more returns become too small to matter. This, firstly, justifies the above second-order expansion. It also means that if we consider covariances of two FC returns we only need to look at first-order terms:

$$\begin{aligned} \mathrm{cov}(\tilde{r}_j^*,\, \tilde{r}_k^*) &\approx \mathrm{cov}(\tilde{r}_j - \tilde{s},\, \tilde{r}_k - \tilde{s}), \\ &= \mathrm{cov}(\tilde{r}_j, \tilde{r}_k) - \mathrm{cov}(\tilde{r}_j, \tilde{s}) - \mathrm{cov}(\tilde{r}_k, \tilde{s}) + \mathrm{var}(\tilde{s}). \end{aligned} \tag{19.42}$$

because all the other terms would lead to products of three or four returns.

One often reads that also inside an expectation only the first-order terms matter, because products of returns are second order of smalls. But this is patently wrong. Indeed, variances and covariances of returns are averages of products of two returns, but this surely does not mean that they can be set equal to zero. Now the expectation of, say, the third term is

$$\mathrm{E}(\tilde{r}_j^*\, \tilde{s}) = \mathrm{E}(\tilde{r}_j^*)\, \mathrm{E}(\tilde{s}) + \mathrm{cov}(\tilde{r}_j^*,\, \tilde{s}). \tag{19.43}$$

If we let the periods over which one observes return become shorter and shorter, all means and all (co)variances shrink roughly in proportion to the time interval Δt, so they preserve the same relative order of magnitude relative to each other. But this means that the product of two means, $\mathrm{E}(\tilde{r}_j^*)\, \mathrm{E}(\tilde{s})$, shrinks to zero much faster than the covariance. That is, the product of two means is second order of smalls but the covariance is not:

$$\mathrm{E}(\tilde{r}_j^*\, \tilde{s}) \approx \mathrm{cov}(\tilde{r}_j,\, \tilde{s}) \text{ , and} \tag{19.44}$$

$$\mathrm{E}(\tilde{s}^2) \approx \mathrm{var}(\tilde{s}); \tag{19.45}$$

Using the above in Equation [19.41], we get the following translated expected return:[16]

$$\mathrm{E}(\tilde{r}_j^*) \approx \mathrm{E}(\tilde{r}_j) - \mathrm{E}(\tilde{s}) - \mathrm{cov}(\tilde{r}_{p^*}^*, \tilde{s}) + \mathrm{var}(\tilde{s}). \tag{19.46}$$

Our results [19.46] and [19.42] for the translated mean and variance imply that the efficiency condition [19.27] translates into the first equation below. We next write that equation for the special case where asset j is the HC risk-free asset, and lastly we subtract:

$$\begin{array}{lllllll} \mathrm{E}(\tilde{r}_j)- & \mathrm{E}(\tilde{s}) & -\mathrm{cov}(\tilde{r}_j, \tilde{s}) & +\mathrm{var}(\tilde{s}) & = \lambda\, [\mathrm{cov}(\tilde{r}_j, \tilde{r}_{p^*}) & -\mathrm{cov}(\tilde{r}_j, \tilde{s}) & -\mathrm{cov}(\tilde{r}_{p^*}, \tilde{s}) + var(\tilde{s})] \\ r- & \mathrm{E}(\tilde{s}) & -0 & +\mathrm{var}(\tilde{s}) & = \lambda\, [0 & -0 & -\mathrm{cov}(\tilde{r}_{p^*}, \tilde{s}) + var(\tilde{s})] \\ \hline \mathrm{E}(\tilde{r}_j) - r & & -\mathrm{cov}(\tilde{r}_j, \tilde{s}) & & = \lambda\, [\mathrm{cov}(\tilde{r}_j, \tilde{r}_{p^*}) & -\mathrm{cov}(\tilde{r}_j, \tilde{s})] & \end{array},$$

which leads to [19.32].

[16]Note, in passing, how we find back our earlier numerical result that covariance between CAD asset return and the CAD/USD exchange rate lowers the expected USD return. We also discover that exchange risk has its impact on the expected return too. So both the covariance and the variance have both 'good' and 'bad' aspects.

Technical Note 19.3 Aggregating the two efficiency conditions.

We want to aggregate, and obtain the world-market return, which is defined as

$$\tilde{r}_w = \frac{W_{ca}\tilde{r}_p + W_{us}\tilde{r}_{p^*}}{W_{ca} + W_{us}} \tag{19.47}$$

with W_{ca} and W_{us} defined as the invested wealths, both measured in CAD, of Canada and the US, respectively. To build this world return into the model we multiply both sides of [19.31] by W_{ca}, and [19.32] by W_{us}. On the right-hand sides of the equations below we have immediately put these factors inside the covariances. Next we sum the two equations, and lastly we divide by total world wealth and use [19.47]:

$$\begin{array}{rll} W_{ca}\mathrm{E}(\tilde{r}_j - r) = & \lambda \operatorname{cov}(\tilde{r}_j,\, W_{ca}\,\tilde{r}_p) & \\ W_{us}\mathrm{E}(\tilde{r}_j - r) = & \lambda \operatorname{cov}(\tilde{r}_j,\, W_{us}\,\tilde{r}_{p^*}) & +W_{us}(1-\lambda)\operatorname{cov}(\tilde{r}_j, s) \\ \hline (W_{ca} + W_{us})\mathrm{E}(\tilde{r}_j - r) = & \lambda \operatorname{cov}(\tilde{r}_j,\, (W_{ca}\,\tilde{r}_p + W_{us}\,\tilde{r}_{p^*})) & +W_{us}(1-\lambda)\operatorname{cov}(\tilde{r}_j, s) \\ & & \\ \Rightarrow \mathrm{E}(\tilde{r}_j - r) = & \lambda \operatorname{cov}(\tilde{r}_j,\, \tilde{r}_w) & +\frac{W_{us}}{W_{ca}+W_{us}}(1-\lambda)\operatorname{cov}(\tilde{r}_j, s). \end{array}$$

For ease of manipulation, in [19.33] we denote $W_{us}/(W_{ca} + W_{us})(1-\lambda) = \kappa$.

Technical Note 19.4 Identifying λ and κ.
Write the equation in matrix form,

$$\mathrm{E}(\tilde{r}_j - r) = [\mathrm{cov}(\tilde{r}_j, r_w)\ ,\ \mathrm{cov}(\tilde{r}_j, \tilde{s})] \begin{bmatrix} \lambda \\ \kappa \end{bmatrix}. \tag{19.48}$$

To identify λ and κ we write this for two benchmarks, the world market portfolio with return r_w and the USD T-bill with return $r^* + \tilde{s}$;

$$\begin{bmatrix} \mathrm{E}(\tilde{r}_w - r) \\ r^* + \mathrm{E}(\tilde{s}) - r \end{bmatrix} = \begin{bmatrix} \mathrm{var}(\tilde{r}_w) & \mathrm{cov}(\tilde{r}_w, \tilde{s}) \\ \mathrm{cov}(\tilde{r}_w, \tilde{s}) & \mathrm{var}(\tilde{s}) \end{bmatrix} \times \begin{bmatrix} \lambda \\ \kappa \end{bmatrix}; \tag{19.49}$$

$$\Rightarrow \begin{bmatrix} \lambda \\ \kappa \end{bmatrix} = \begin{bmatrix} \mathrm{var}(\tilde{r}_w) & \mathrm{cov}(\tilde{r}_w, \tilde{s}) \\ \mathrm{cov}(\tilde{r}_w, \tilde{s}) & \mathrm{var}(\tilde{s}) \end{bmatrix}^{-1} \begin{bmatrix} \mathrm{E}(\tilde{r}_w - r) \\ r^* + \mathrm{E}(\tilde{s}) - r \end{bmatrix}. \tag{19.50}$$

This can be substituted back into [19.48]. Now the covariance matrix of ($\tilde{r}_w$, $\tilde{s}$) premultiplied by the vector of covariances of r_j with these same variables ($\tilde{r}_w$, $\tilde{s}$) is the row vector of multiple regression coefficients of r_j onto ($\tilde{r}_w$, $\tilde{s}$)—a generalisation of $b = \mathrm{cov}(\tilde{y}, \tilde{x}) \times \mathrm{var}(\tilde{x})^{-1}$ in $\tilde{y} = a + b\tilde{x} + \tilde{e}$:

$$\begin{aligned} \mathrm{E}(\tilde{r}_j - r) &= [\mathrm{cov}(\tilde{r}_j, r_w)\ ,\ \mathrm{cov}(\tilde{r}_j, \tilde{s})] \begin{bmatrix} \mathrm{var}(\tilde{r}_w) & \mathrm{cov}(\tilde{r}_w, \tilde{s}) \\ \mathrm{cov}(\tilde{r}_w, \tilde{s}) & \mathrm{var}(\tilde{s}) \end{bmatrix}^{-1} \begin{bmatrix} \mathrm{E}(\tilde{r}_w - r) \\ r^* + \mathrm{E}(\tilde{s}) - r \end{bmatrix} \\ &= [\beta_{j,w;s}\ ,\ \gamma_{j,s;w}] \begin{bmatrix} \mathrm{E}(\tilde{r}_w - r) \\ r^* + \mathrm{E}(\tilde{s}) - r \end{bmatrix}. \end{aligned} \tag{19.51}$$

Technical Note 19.5 The best-replication reading of the i-CAPM

The claim can be shown as follows. In the first line, we write the return on a general portfolio with weights x_w and x_s for the world market and the foreign T-bill, and in the second line we group terms in x_w and x_s:

$$\begin{aligned} \mathrm{E}(\tilde{r}_{\text{j's replication}}) &= x_w \mathrm{E}(\tilde{r}_w) + x_s(r^* + \tilde{s}) + (1 - x_w - x_s)r && (19.52)\\ &= r + x_w \mathrm{E}(\tilde{r}_w - r) + x_s(r^* + \mathrm{E}(\tilde{s}) - r). && (19.53) \end{aligned}$$

For best replication, we have to set $x_w = \beta_{j,w;s}$ and $x_s = \gamma_{j,s;w}$. Thus,

$$\mathrm{E}(\tilde{r}_{\text{j's replication}} - r) = \beta_{j,w;s} \mathrm{E}(\tilde{r}_w - r) + \gamma_{j,s;w}(r^* + \mathrm{E}(\tilde{s}) - r). \quad (19.54)$$

19.6 Test Your Understanding: basics of the CAPM

19.6.1 Quiz Questions

True-False Questions

1. The risk of a portfolio is measured by the standard deviation of its return.
2. The risk of an asset is measured by the standard deviation of its return.
3. Each asset's contribution to the total risk of a portfolio is measured by the asset's contribution to the total return on the portfolio.
4. A risk-averse investor always prefers the highest possible return for a given level of risk or the lowest risk for a given level of expected return.
5. The means and standard deviations of all optimal portfolios selected from a risk-free asset and a set of risky assets are found on the line that originates at r_0 and is tangent to the efficient portfolio of risky assets.
6. Relative risk aversion shows the price in currency units of a given amount of risk.
7. Relative risk aversion varies from asset to asset because some assets are riskier than others.
8. Portfolio theory assumes that all investors are equally risk averse.

Multiple-Choice Questions

1. When using portfolio theory, we must make a number of assumptions. Which of the following assumptions are made? Which are not?
   (a) The rates of inflation at home and abroad are equal.
   (b) There are no information or transactions costs.
   (c) There are no taxes.
   (d) Investors want to know the distribution of wealth at the end of the period.
   (e) Investors care about the future expected return on their portfolio and the variability of this return.

19.6.2 Applications

1. The Country Prince Rupert's Land (PRL) has two companies, Hudson Bay Company (HBC) and Boston Tea Traders (BTT). In equilibrium, the returns of these two companies have the following distributions:

	Expected	Covariances	
	excess return	HBC	BTT
HBC	0.11	0.04	0.01
BTT	0.08	0.01	0.02

(a) Vary the weight of HBC from 0 to 1 by increments of 0.1, and compute how the portfolio covariance risks of HBC and BTT change as a function of the weights x_{HBC} and $x_{BTT} = 1 - x_{HBC}$.

(b) Find the optimal weights of x_{HBC} and $x_{BTT} = 1 - x_{HBC}$ and the average risk aversion.

(c) If the total value of the PRL stock market portfolio is 1,000, what is the value of HBC and BTT?

2. Consider the following covariance matrix and expected return vector for assets 1, 2, and 3:

$$V = \begin{bmatrix} 0.0100 & 0.0020 & 0.0010 \\ 0.0020 & 0.0025 & 0.0030 \\ 0.0010 & 0.0030 & 0.0100 \end{bmatrix} \quad E(\tilde{r}_j) = \begin{bmatrix} 0.0330 \\ 0.0195 \\ 0.0250 \end{bmatrix}$$

(a) Compute the expected return on a portfolio with weights for assets $j = 0, \ldots, 3$ equal to $[0.2, 0.4, 0.2, 0.2]$, when the T-bill (asset 0) yields a return of 1 percent. Do so directly, and then via the excess returns.

(b) Compute the variance of the same portfolio.

(c) Compute the covariance of the return on each asset with the total portfolio return, and verify that it is a weighted covariance.

(d) Is the above portfolio efficient?

(e) Are the following portfolios efficient?
• weights $(0.7, 0.1, 0.1, 0.1)$ for assets $j = 0, \ldots, 3$
• weights $(0.6, 0.2, 0.1, 0.1)$ for assets $j = 0, \ldots, 3$

(f) What is the portfolio held by an investor with risk-aversion measure $\lambda = 2.5$?

(g) Assume that there are no "outside" bills, that is, all risk-free lending and borrowing is among investors. Therefore the average investor holds only risky assets. What is the portfolio composition? What is the average investor's risk-aversion measure λ?

19.7 Test Your Understanding: iCAPM

19.7.1 Quiz Questions

True-False Questions

1. The entire NPV analysis can be conducted in terms of the host (foreign) currency if money markets and exchange markets are fully integrated with the home market.

2. The entire NPV analysis can be conducted in terms of the host currency if money markets, stock markets, and exchange markets are fully integrated with the home market.

3. Forward rates can be used as the risk-adjusted expected future spot rates to translate the host-currency cash flows into the home currency. The home-currency cash flows can then be discounted at the appropriate home-currency discount rate if money markets and exchange markets are fully integrated with the home market.

4. Regardless of the degree of market integration, the host-currency expected cash flows can always be translated into the home currency (by multiplying them by the expected spot rate), and then discounted at the home-currency discount rate.

5. Regardless of the degree of market integration, the host-currency expected cash flows can always be translated into expected cash flows expressed in home currency. The home-currency cash flows can then be discounted at the home-currency discount rate that takes into account all risks.

6. If you use the forward rate as the risk-adjusted expected spot rate, there is no need to worry about the dependence between the exchange rate and the host-currency cash flows.

7. If markets are integrated and you translate at the forward rate, the cost of capital need not include a risk premium for exchange rate exposure.

8. If markets are integrated and you translate at the forward rate, the cost of capital need not include a risk premium for exposure to any currency.

9. If you discount expected cash flows that are already expressed in home currency, the cost of capital should include a risk premium for exposure to the host-currency exchange rate.

10. If you discount expected cash flows that are already expressed in home currency, the cost of capital should include a risk premium for exposure to all relevant exchange rates.

11. If you translate at the forward rate, you can entirely omit exchange rate expectations from the NPV procedure.

12. Exchange rate risk premia are sizeable. In fact, they are about as large as the (world) market risk premium.

13. A highly risk-averse investor will only accept variance risk if he or she is fully certain to be compensated for this risk.

14. A highly risk-averse investor will never select a high-variance portfolio.

15. A risk-averse investor will select a high-variance portfolio only if the expected excess return is sufficiently high.

16. A risk-averse investor will select a low-return portfolio only if the variance is sufficiently low.

17. A particularly risk-averse investor will always select a low-return portfolio. This is because low return means low risk, and because the investor does not want to bear a lot of risk.

For the next set of questions, assume that access to money markets and exchange markets is unrestricted and the host-currency cash flow is risk free. Are the following statements true or false?

18. You can translate at the expected spot rate and discount at a risk-adjusted home-currency cost of capital.

19. You can translate at the forward rate, and discount at a home-currency rate that takes into account exchange risk.

20. You can translate at the forward rate, and discount at the risk-free home-currency rate.

21. You can discount the host-currency cash flows at the foreign risk-free rate, and then translate the result at the current spot exchange rate.

22. You can discount the host-currency cash flows at the foreign risk-free rate, and then translate the result at the expected future spot exchange rate.

23. You can discount the host-currency cash flows at the foreign risk-free rate, and then translate the result at the forward exchange rate.

24. If access to forward markets or foreign and domestic money markets is restricted, then the true value is always overstated if the foreign currency cash flow is translated at the forward exchange rate and then discounted at the domestic risk-free rate.

Additional Quiz Questions

1. Suppose that you observe an efficient portfolio. There are two methods with which you can infer the degree of risk aversion of the investor that selects this particular portfolio. What are these two methods?

2. What's wrong with the following statement: "The CAPM says that the expected return on a given stock j is equal to the best possible replication that one can obtain using the risk-free assets and the set of all risky assets (other than stock j)."

3. Below, we reproduce some equations from the derivation of the CAPM. Equation [20.1] is the efficiency criterion. Equation [19.62] is the CAPM. Explain the equations.

$$\frac{\mathrm{E}(\tilde{r}_j - r)}{\mathrm{cov}(\tilde{r}_j - \tilde{r}_m)} = \theta, \tag{19.55}$$

for all risky assets j=1,...,N.

$$\begin{aligned} \mathrm{E}(\tilde{r}_j - r) &= \theta\ \mathrm{cov}(\tilde{r}_j, \tilde{r}_m), & (19.56)\\ &= [\theta\ \mathrm{var}(\tilde{r}_m)]\frac{\mathrm{cov}(\tilde{r}_j, \tilde{r}_m)}{\mathrm{var}(\tilde{r}_m)}, & (19.57)\\ &= [\theta\ \mathrm{var}(\tilde{r}_m)]\beta_j, & (19.58)\\ \sum_{j=1}^{N} x_j \mathrm{E}(\tilde{r}_j - r) &= \theta\ , & (19.59)\\ \sum_{j=1}^{N} x_j \mathrm{cov}(\tilde{r}_j, \tilde{r}_m) &= \theta\ \mathrm{cov}(\sum_{j=1}^{N} x_j \tilde{r}_j, \tilde{r}_m), & (19.60)\\ &= \theta\ \mathrm{cov}(\tilde{r}_m, \tilde{r}_m), & (19.61) \end{aligned}$$

$$\mathrm{E}(\tilde{r})_j - r = \beta_j[\mathrm{E}(\tilde{r}_m) - r]. \tag{19.62}$$

4. Suppose that investors from a country have access to a large set of foreign stocks, and that foreign investors can also buy stocks in that country. Which of the following statements is (are) correct?

   (a) The single-market CAPM, where the market portfolio is measured by the index of all stocks issued by local companies, does not hold.

   (b) The single-market CAPM, where the market portfolio is measured by the index of all stocks held by local investors, does not hold.

   (c) The single-market CAPM, where the market portfolio is measured by the index of all stocks held by local investors, is formally correct but not fit for practical use, because the correct index is not readily observable.

   (d) The single-market CAPM, where the market portfolio measured by the index of all stocks worldwide, is correct provided that there is a unified world market for all stocks.

(e) The single-market CAPM, where the market portfolio is measured by the index of all stocks worldwide, is correct provided that there is no (real) exchange risk.

19.7.2 Applications

1. Suppose that you have the following data:
   Asset 0 is the (domestic) risk-free asset, and asset weights in a portfolio are denoted as x_j, where j = 0,.. ., 2. Which of the following portfolios is efficient, and if the portfolio is efficient, what is the investor's degree of risk aversion?

   (a) $x_0 = 0$, $x_1 = 0.4$, $x_2 = 0.6$
   (b) $x_0 = 0$, $x_1 = 0.6$, $x_2 = 0.4$
   (c) $x_0 = 0$, $x_1 = 0.5$, $x_2 = 0.5$
   (d) $x_0 = 0.2$, $x_1 = 0.4$, $x_2 = 0.4$
   (e) $x_0 = 0.5$, $x_1 = 0.25$, $x_2 = 0.25$
   (f) $x_0 = –1$, $x_1 = 1$, $x_2 = 1$
   (g) $x_0 = 1$, $x_1 = 0$, $x_2 = 0$
   (h) $x_0 = 2$, $x_1 = –0.5$, $x_2 = –0.5$

2. Suppose that the capital markets of the following three countries are well integrated: North America (with the dollar), Europe (with the EUR), and Japan (with the yen). Suppose that you choose the yen as the home currency.

   (a) Why does the average investor care about the JPY/USD and JPY/EUR exchange rates (beside how it relates to how his or her wealth is measured in JPY)?
   (b) What moments are needed in a mean-and-(co)variance framework, to summarize the joint distribution of asset returns? Which of these are affected by the portfolio choice?

3. Suppose that your assistant has run a market-model regression for a company that produces sophisticated drilling machines, and finds the following results (t-statistic in parentheses):

$$\tilde{r}_j = \alpha + \beta\tilde{r}_m + \gamma s + \tilde{e}_j,$$
$$\tilde{r}_j = \underset{(0.52)}{0.002} + \underset{(1.25)}{0.56}\tilde{r}_m + \underset{(2.06)}{4.25}\tilde{s} + \tilde{e}_j.$$

   Your assistant remarks that, as the estimated beta is insignificant, the true beta is zero. The exposure, in contrast, is significant, and must be equal to the estimated coefficient. How do you react?

4. Suppose that the world beta for a German stock (in euro) equals 1.5, and its exposures to the dollar, the yen, and the pound are 0.3, 0.2, and 0.1, respectively.

   (a) What is the best replicating portfolio if you can invest in a world-market index fund, as well as in dollars, yens, pounds, and euros?

   (b) What additional information is needed to identify the cost of capital?

5. Suppose that there are two countries, the US (which is the foreign country) and Canada. The exposure of the company XUS, in terms of USD, is estimated as follows:

$$\tilde{r}^*_{XUS} = 0.12 + 0.30\tilde{s}_{USD/CAD} + \tilde{\varepsilon}.$$

   What is the company's exposure in terms of CAD?

