

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 collateralizing.

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 counterparty, 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:

$$\begin{array}{rcl}
 \text{- time 1: pay } & \underbrace{(40 - 38)/1.02 \times 1.02}_{1.961 \text{ borrowed}} & = 40 - 38 = & 2 \\
 \text{- time 2: pay } & \underbrace{(38 - 36)/1.01 \times 1.01}_{1.98 \text{ borrowed}} & = 38 - 36 = & 2 \\
 \text{- time 3: pay} & & & \underline{36} \\
 \text{- total: pay} & & & 40
 \end{array}$$

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: $1.980 \times 1.01 =$ 2m - time 1: $1.961 \times 1.02 =$ 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 counterparty. 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 simplification 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 payments 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 $f_{t_0,T} = \text{USD}/\text{EUR } 0.96$. 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, $1\text{m} \times (0.96 - 0.89) = \text{USD } 70,000$ 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 $1\text{m} \times 0.001 = \text{USD } 1,000$. 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, $1\text{m} \times (F_{t_0,T} - S_T) = 1\text{m} \times (0.96 - 0.889) = \text{USD } 71,000$.

- 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

¹The 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 Jérôme 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 fictitious 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

²LIFFE: 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, intransparent 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.

³DTB: 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 $12.5\text{m} \times (0.0046/100) = \text{USD } 575$ 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 $F_{0,2} = f_{0,2}$

$F_{1,2}$	HC cash flow: futures		HC cash flow: forward		difference	
	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 $F_{0,2} = \text{USD } 100$. 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 $F_{0,2} = f_{0,2}$

state	case 1			case 2			case 3		
	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:

$$\text{basis} := f_{t,T} - S_t. \quad (6.1)$$

⁴If 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, unrealistically, a constant SEK/EUR cross rate, the hedger would have to sell eight contracts to approximately hedge the SEK 9,000,000: $8 \times 1,162,500 = 9.3\text{m}$. 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

⁵Beta 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 choosing 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

$$\text{Cash flow at time } T_1 = \tilde{S}_{T_1}^{(e)} - \beta \times (\tilde{f}_{T_1, T_2}^{(h)} - f_{t, T_2}^{(h)}). \quad (6.2)$$

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:

$$\begin{aligned} \beta &= \text{the slope coefficient from } \tilde{S}_{T_1}^{(e)} = \alpha + \beta\tilde{f}_{T_1, T_2}^{(h)} + \tilde{\epsilon}, \\ &= \frac{\text{cov}(\tilde{f}_{T_1, T_2}^{(h)}, \tilde{S}_{T_1}^{(e)})}{\text{var}(\tilde{f}_{T_1, T_2}^{(h)})}. \end{aligned} \quad (6.3)$$

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 $\tilde{f}_{T_1, T_2}^{(e)} = S_{T_1}^{(e)}$: 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

$$\beta = \frac{\text{cov}(\tilde{S}_{T_1}^{(e)}, \tilde{S}_{T_1}^{(h)})}{\text{var}(\tilde{S}_{T_1}^{(h)})}, \quad (6.4)$$

$$= \text{the slope coefficient in } \tilde{S}_{T_1}^{(e)} = \alpha + \beta \tilde{S}_{T_1}^{(h)} + \tilde{\epsilon}. \quad (6.5)$$

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:

$$\text{regress } \Delta S_t^{(e)} = \alpha' + \beta \Delta S_t^{(h)} + \tilde{\epsilon}'_t, \quad (6.6)$$

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 the distribution of ΔS is constant. So if we use s as shorthand for $\Delta S/S$, we would use the following equation:

$$\tilde{s}_t^{(e)} = \alpha'' + \gamma \tilde{s}_t^{(h)} + \tilde{\epsilon}_t'' \quad (6.7)$$

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

$$\text{if } \tilde{s}_t^{(e)} = \alpha'' + \gamma \tilde{s}_t^{(h)} + \tilde{\epsilon}_t'' \text{ then } \beta = \gamma \frac{S_t^{(e)}}{S_t^{(h)}} \quad (6.8)$$

$$= \gamma S_t^{(cross)}, \quad (6.9)$$

where $S^{(cross)}$ is the cross rate. In our example,

$$\frac{S_t^{(e)}}{S_t^{(h)}} \text{ has dimension } \frac{\text{USD/SEK}}{\text{USD/EUR}} = \text{EUR/SEK}, \quad (6.10)$$

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:

$$\text{Rule of thumb for cross hedge: } \gamma = 1 \text{ so } \beta = S_t^{(h/e)}. \quad (6.11)$$

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 one 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,

⁶In 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 indis-

tinguishable from forwards, we know that

$$\tilde{f}_{T_1, T_2}^{(e)} \approx \tilde{S}_{T_1}^{(e)} \frac{1 + \tilde{r}_{T_1, T_2}}{1 + \tilde{r}_{T_1, T_2}^{(e)}}, \quad (6.12)$$

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

$$\tilde{S}_{T_1}^{(e)} = \frac{1 + \tilde{r}_{T_1, T_2}^{(e)}}{1 + \tilde{r}_{T_1, T_2}} \tilde{f}_{T_1, T_2}^{(e)}. \quad (6.13)$$

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

$$\tilde{S}_{T_1}^{(e)} = 1.005 \tilde{f}_{T_1, T_2}^{(e)}, \quad (6.14)$$

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.

⁷This 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

$$\text{hedge ratio} = \frac{n_e}{n_h} \beta, \quad (6.15)$$

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:

$$\Delta S_{[\text{USD}/\text{SEK}]} = 0.003 + 0.105 \Delta f_{[\text{USD}/\text{EUR}]}. \quad (6.16)$$

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

$$\text{hedge ratio} = \frac{2,170,000}{125,000} \times 0.105 = 1.822, \quad (6.17)$$

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 effects. 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 Options 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 (São 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_{t,T_1,T_2}^f . 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:

$$V_t = \frac{1}{1 + r_{t,T_2}}, \quad (6.18)$$

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:

$$\begin{aligned}
 V_{t,T_1,T_2}^f &= V_t (1 + r_{t,T_1}) \\
 &= \frac{1 + r_{t,T_1}}{1 + r_{t,T_2}}, && \text{see [6.18],} \\
 &= \frac{1}{1 + r_{t,T_1,T_2}^f}, && \text{see [4.48].}
 \end{aligned} \tag{6.19}$$

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:

$$V_{t,T_1,T_2}^f = 97.087 \times 1.0195 = 98.981\%. \tag{6.21}$$

Alternatively, we can compute the six-month forward price on a nine-month T-bill via the forward rate of return:

$$1 + r_{t,T_1,T_2}^f = \frac{1.03}{1.0195} = 1.010,299 \Rightarrow V_{t,T_1,T_2}^f = \frac{1}{1.010,299} = 98,981\%. \tag{6.22}$$

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

$$V_{t,T_1,T_2}^f = \frac{1}{1 + (1/4)0.041} = 98.985,300 \approx 98,99. \quad (6.24)$$

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 $1m \times 0.0001 =$ 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 $1\text{m}/(1 + (1/4) 0.04) = 990,099.01$ to $1\text{m}/(1 + (1/4) 0.039) = 990,344.14$, 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, FRAs 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 $E_t(\tilde{y}) = E_t(\tilde{x})$. 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 $E(\tilde{y}) = a + bE(\tilde{x}) + 0$ 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,

$$E_t(\tilde{s}_T^{sek}) = E_t(\tilde{s}_T^{eur}) \text{ does not imply } E_t(\tilde{s}_T^{sek} | s_T^{eur}) = s_T^{eur}, \quad (6.26)$$

even though the reverse statement does hold:

$$E_t(\tilde{s}_T^{sek} | s_T^{eur}) = s_T^{eur} \text{ does imply } E_t(\tilde{s}_T^{sek}) = E_t(\tilde{s}_T^{eur}). \quad (6.27)$$

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 $E_t(\tilde{s}_T^{sek} | s_T^{eur}) = s_T^{eur}$ (the crown rises as much as the euro)
- b. if there's euro news, then $E_t(\tilde{s}_T^{sek} | s_T^{eur}) = 0$ (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.

	Open	High	Low	Settle	Change	Lifetime		Open
						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.
- What are the daily cash flows from marking to market?
 - What is the total cash flow from marking to market (ignoring discounting)?
 - 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, % <i>p.a.</i>	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.
- Your assistant runs a bunch of regressions:
 - $\Delta S[\text{EUR}/\text{CAD}] = \alpha_1 + \beta_1 \Delta f[\text{USD}/\text{EUR}]$
 - $\Delta S[\text{EUR}/\text{CAD}] = \alpha_2 + \beta_2 \Delta f[\text{EUR}/\text{USD}]$
 - $\Delta S[\text{CAD}/\text{EUR}] = \alpha_3 + \beta_3 \Delta f[\text{EUR}/\text{USD}]$
 - $\Delta S[\text{CAD}/\text{EUR}] = \alpha_4 + \beta_4 \Delta f[\text{USD}/\text{EUR}]$
 Which regression is relevant to you?
 - 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 $R^2 = 0.59$:

$$\Delta S_{[\text{EUR}/\text{NZD}]} = a + \underset{(1.57)}{0.15} \Delta f_{[\text{EUR}/\text{USD}]} + \underset{(17.2)}{0.7} \Delta f_{[\text{EUR}/\text{AUD}]}$$

- 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?

