

Chapter

7

Speculation and Risk in the Foreign Exchange Market

Japanese investors, like Mrs. Watanabe in Chapter 2, have faced perennially low Japanese yen interest rates for years. They consequently have found high-yielding bonds denominated in Australian and New Zealand dollars quite attractive. More recently, retail aggregator accounts have been introduced that allow private Japanese investors to speculate in foreign exchange markets using forward contracts. A 2010 Bank for International Settlements (BIS) study by Michael King and Dagfinn Rime estimates that Japanese retail investors trade over \$20 billion a day in foreign exchange markets.

This chapter examines how investors quantify expected returns and risks associated with speculative foreign exchange investments. If an investor chooses not to hedge (or “cover”) the exchange risk on a foreign money market investment, the return is uncertain and will be high if the foreign currency appreciates or low if the foreign currency depreciates. Our discussion of uncovered investments in the foreign money market uses some basic statistical methods that are commonly used to explain empirical evidence about investment returns in all asset markets. The Appendix to Chapter 3 and Appendix 7.3 in this chapter provide the necessary background.

7.1 SPECULATING IN THE FOREIGN EXCHANGE MARKET

Uncovered Foreign Money Market Investments

In Chapter 6, we examined covered foreign money market investments and found that if interest rate parity is satisfied, the domestic currency rate of return from investing in a foreign money market and covering the foreign exchange risk is the domestic currency interest rate. What happens if an investor does not cover the foreign exchange risk? Let’s look at an example.

Example 7.1 Kevin Anthony's Uncovered Pound Investment

Recall the situation in Example 6.2 in which Kevin Anthony, a portfolio manager, was considering several ways to invest \$10,000,000 for 1 year. The data are as follows:

USD interest rate: 8.0% per annum (p.a.)
 GBP interest rate: 12.0% p.a.
 Spot exchange rate: \$1.60/£

Remember that if Kevin invests in the USD-denominated asset at 8%, after 1 year he will have $\$10,000,000 \times 1.08 = \$10,800,000$.

Suppose Kevin invests his \$10,000,000 in the pound money market, but he decides not to hedge the foreign exchange risk. As before, we can calculate his dollar return in three steps.

Step 1. Convert dollars into pounds in the spot market. The \$10,000,000 will buy

$$\frac{\$10,000,000}{\$1.60/\text{£}} = \text{£}6,250,000$$

at the current spot exchange rate. This is Kevin's pound principal.

Step 2. Calculate pound-denominated interest plus principal. Kevin can invest his pound principal at 12%, yielding a return in 1 year of

$$\text{£}6,250,000 \times 1.12 = \text{£}7,000,000$$

Step 3. Sell the pound principal plus interest at the spot exchange rate in 1 year:

$$\text{Dollar proceeds in 1 year} = \text{£}7,000,000 \times S(t+1, \$/\text{£})$$

By choosing not to hedge the foreign exchange risk, the dollars Kevin receives from his investment in the pound money market are determined by the value of the future exchange rate.

Let's denote the \$/£ current spot rate by $S(t)$ and the future spot rate by $S(t+1)$. Following the three steps in Example 7.1, the dollar return from investing 1 dollar in a pound money market investment, $r(t+1)$, is

$$r(t+1) = \frac{1}{S(t)} \times [1 + i(\text{£})] \times S(t+1) \quad (7.1)$$

where $i(\text{£})$ denotes the pound interest rate. In Example 7.1, we obtain

$$r(t+1) = \frac{1}{\$1.60/\text{£}} \times 1.12 \times S(t+1) = 0.7 \times S(t+1).$$

Notice that $0.7 = \frac{\text{£}7,000,000}{\$10,000,000}$ is the ratio of the amount of future pounds Kevin will have

to the amount of dollars he invests today. The return on Kevin's investment is risky because the value of the future exchange rate is not known today. Kevin might also be interested in the excess return to this investment, denoted $\text{exr}(t+1)$ —that is, the return over and above what he could earn risk free domestically. The excess return (exr) is

$$\begin{aligned} \text{exr}(t+1) &= \frac{S(t+1)}{S(t)} \times [1 + i(\text{£})] - [1 + i(\$)] \\ &= S(t+1) \times 0.7 - 1.08 \end{aligned} \quad (7.2)$$

where $i(\$)$ is the dollar interest rate.

Speculating with Forward Contracts

The Break-Even Spot Rate

The future exchange rate for which Kevin breaks even between the pound and the domestic money market investments is the exchange rate, S^{BE} , that sets Equation (7.2) equal to zero:

$$S^{\text{BE}} = S(t) \times \frac{[1 + i(\$)]}{[1 + i(\pounds)]} \quad (7.3)$$

Hence, Kevin's break-even rate is $S^{\text{BE}} = 1.08/0.7 = \$1.5429/\pounds$.

From Chapter 6, recognize that Equation (7.3) is the formula for the forward rate! Consequently, if the foreign currency appreciates such that the future exchange rate is above the forward rate, Kevin makes a positive excess return, but if the future exchange rate is less than the forward rate, Kevin has a negative excess return. Therefore, it is not surprising that Kevin can also speculate on the direction of the pound exchange rate using forward contracts.

Comparing Forward Market and Foreign Money Market Investments

Forward contracts are pure bets—that is, no money changes hands when a forward contract is made. To make this forward contracting situation more concrete, let Mr. Buy represent the person who buys pounds forward with dollars from Ms. Sell, who represents the person who sells pounds forward for dollars. Mr. Buy will pay $F(t)$ dollars in 1 year for every pound he buys forward, and he will sell each pound in the future spot market for dollars at $S(t+1)$. Ms. Sell, on the other hand, will buy her pounds in the future spot market at a dollar price of $S(t+1)$, and she will sell each pound to Mr. Buy for $F(t)$. Therefore, on a per-pound basis, the dollar profits and losses are as follows:

$$\text{Mr. Buy's dollar profit or loss} = S(t+1) - F(t)$$

$$\text{Ms. Sell's dollar profit or loss} = F(t) - S(t+1)$$

These dollar profits and losses are graphed in Exhibit 7.1 as a function of $S(t+1)$. Notice that the dollar profit of the person buying foreign currency forward is the dollar loss of the person selling foreign currency forward, and vice versa.

How does this **forward market investment** compare with Kevin Anthony's pound foreign money market investment? Because Kevin invests in the pound money market, the relevant comparison is with Mr. Buy's purchase of pounds in the forward market. We first express Mr. Buy's profits on a per-dollar basis by dividing by $S(t)$:

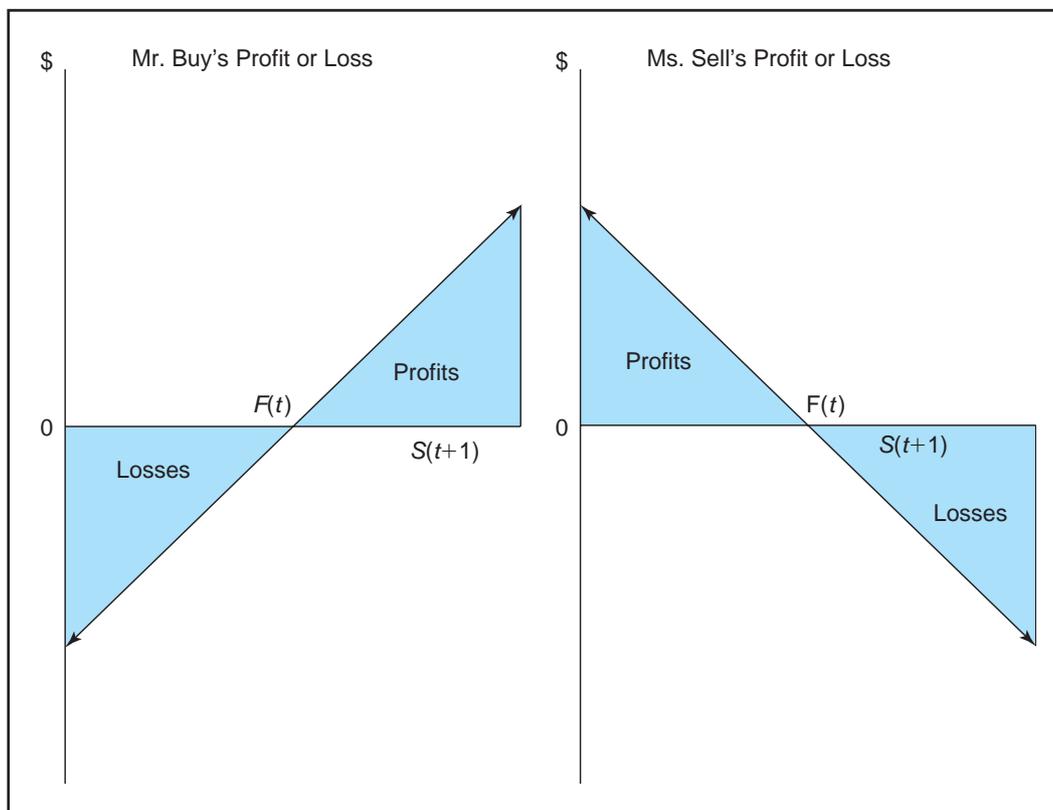
$$\text{Forward Market return (per dollar)} = \text{fmr}(t+1) = \frac{S(t+1) - F(t)}{S(t)} \quad (7.4)$$

where we define the **forward market return** (per dollar) in Equation (7.4) as $\text{fmr}(t+1)$. Because the excess return can be viewed as the return on a strategy in which Kevin borrows dollars in the domestic money market and invests them in the pound money market, it is analogous to a forward contract in which no money changes hands up front. Clearly, the two returns must be closely related, as both investments are exposed to changes in the value of the pound. In fact,

$$\text{fmr}(t+1) \times [1 + i(\pounds)] = \frac{S(t+1)}{S(t)} [1 + i(\pounds)] - [1 + i(\$)]$$

Intuitively, because the forward contract sells £1 in the future, but Kevin's strategy invests pounds today, we must make a future value adjustment. We must scale up the forward market return by $[1 + i(\pounds)]$ to compare it to a money market investment because 1 pound today is worth $[1 + i(\pounds)]$ pounds in the future. Mathematically, you can verify this relation by replacing $F(t)$ in the expression for $\text{fmr}(t)$ by its value in terms of the spot exchange rate and interest rates predicted by covered interest rate parity [see Equation (7.3)].

Exhibit 7.1 Profits and Losses from Forward Market Speculation



Currency Speculation and Profits and Losses

The uncertainty about future exchange rates makes currency speculation risky. We now show how to characterize expected losses and profits on speculative currency investments.

Quantifying Expected Losses and Profits

To quantify our uncertainty about future returns, we use conditional probability distributions as in Chapter 3. Recall that we view today as being time t , and remember that the *conditional probability distribution* of the spot exchange rate for some time in the future, as in Exhibit 3.1, describes the conditional probabilities associated with all the possible exchange rates that may occur at that time *conditioned on* all the information that is available today. The collection of all information that is used to predict the future value of an economic variable is typically called an **information set**. Also, recall that we refer to the expected value (the mean) of this probability distribution as the *conditional expectation of the future exchange rate*. We denote the conditional expectation at time t of the future spot exchange rate of dollars per pound at time $t+1$, for instance, 1 year from now, as $E_t[S(t+1, \$/\pounds)]$.

In Chapter 3, we argued that the distribution of exchange rate changes is relatively well described by a normal (that is, a bell-shaped) distribution, at least for exchange rates between the currencies of developed countries. As we will argue later in this chapter, there are times when conditional distributions of future exchange rates are fat tailed and skewed. For now, though, we'll stick to the normal distribution because it often works well. Hence, in addition to the mean of the conditional distribution of the future spot exchange rate, we must also specify its standard deviation. Now we are ready to quantify the probability of losses and gains. Let's illustrate by revisiting Kevin Anthony's example.

Example 7.2 Kevin Anthony's Probability of Loss

Suppose Kevin expects the pound to depreciate relative to the dollar by 3.57% over the next year. Then, the conditional expectation of his future spot rate in 1 year is

$$\$1.60/\text{£} \times (1 - 0.0357) = \$1.5429/\text{£}$$

which makes the conditional expectation of his uncertain dollar return equal to

$$£7,000,000 \times \$1.5429/\text{£} = \$10,800,300$$

This return is essentially the same as the return from his dollar investment because $\$1.5429/\text{£}$ is the break-even future exchange rate (S^{BE}) that equalizes the returns on dollar and pound investments.¹

Suppose Kevin thinks that the rate of appreciation of the pound relative to the dollar is normally distributed. From the symmetry of the normal distribution, he knows that there is a 50% probability that he will do better than the dollar investment and there is a 50% probability that he will do worse.

Kevin might also be interested in knowing the probability that he will lose some of his dollar principal. At what future value of the spot exchange rate $S(t+1, \$/\text{£})$ will Kevin just get his $\$10,000,000$ principal back? This value—let's call it \hat{S} —satisfies

$$(\text{£}7,000,000) \times \hat{S} = \$10,000,000$$

from which we find

$$\hat{S} = \frac{\$10,000,000}{\text{£}7,000,000} = \$1.4286/\text{£}$$

Kevin can calculate the probability that the future exchange rate will be lower than $\$1.4286/\text{£}$. To perform such a calculation, he needs to determine the standard deviation of the payoff on his investment. Suppose he thinks that the standard deviation of the rate of appreciation of the pound relative to the dollar over the next year is 10%. Because 10% of $\$1.60/\text{£}$ is $\$0.16/\text{£}$, the standard deviation of the conditional distribution of the future spot exchange rate is $\$0.16/\text{£}$ (see Chapter 3). He can calculate the probability of losing money by creating a **standard normal random variable**. A standard normal random variable has a mean of 0 and a standard deviation of 1, which we denote with $N(0, 1)$, and we can calculate it by subtracting the mean of the future spot rate and dividing by the standard deviation. Thus,

$$\frac{S(t+1, \$/\text{£}) - \$1.5429/\text{£}}{\$0.16/\text{£}}$$

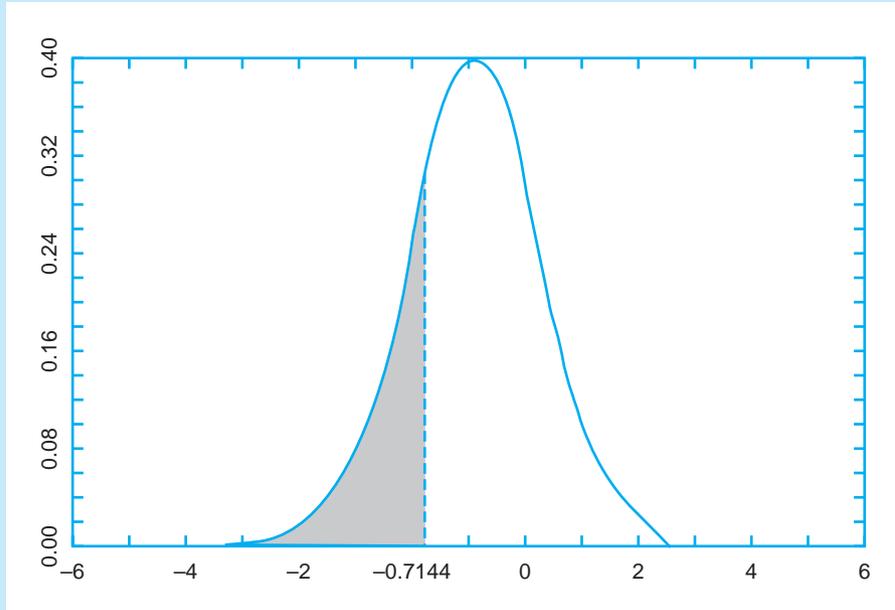
has a mean of 0 and a standard deviation of 1. We graph such a standard normal distribution in Exhibit 7.2. Then, the value of the standard normal variable associated with a zero rate of return is

$$\frac{\$1.4286/\text{£} - \$1.5429/\text{£}}{\$0.16/\text{£}} = -0.7144$$

From the probability distribution of a standard normal, we find that there is a 23.75% probability that a $N(0, 1)$ variable will be less than -0.7144 , or equivalently that $S(t+1, \$/\text{£})$ will be less than $\$1.4286/\text{£}$. In the graph in Exhibit 7.2, the area below the curve to the left of -0.7144 is 23.75% (the total area sums to 1). Hence, 23.75% is the chance that Kevin will actually lose some of his dollar principal over the course of the next year.

¹The $\$300$ difference is due to the rounding of the exchange rate to the fourth digit.

Exhibit 7.2 Standard Normal Distribution



Notes: The horizontal axis represents possible values for a standard normally distributed variable (say, x). The vertical axis represents the value of the normal distribution function (say, y) for each x . In fact,

$y = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2}$, where e is 2.71828. The area below -0.7144 represents 23.75% of the total area, which sums to 1.

Lessons from History: The Variability of Currency Changes and Forward Market Returns

At this point, one can think of the conditional probability distribution as reflecting the subjective beliefs of an individual investor, an importer or an exporter, about the uncertain future exchange rate. The next section discusses theories that determine a value for the conditional mean of the distribution. Here we review historical data to inform us about the width of the distribution. Kevin used 10% for the rate of appreciation of the pound versus the dollar. If the true number were larger, the conditional distribution for the future exchange rate would be more dispersed, and the probability that he would lose some of his principal would be larger than 23.75%.

Exhibit 7.3 shows the standard deviations of percentage changes in exchange rates and forward market returns for three exchange rates versus the U.S. dollar and the corresponding non-dollar cross rates calculated with over 30 years of actual data. The three currencies are the euro (using data on the Deutsche mark before 1999), the British pound, and the Japanese yen. Note that the annualized volatilities of percentage changes in the exchange rate reported in column 1 are indeed around 10% (somewhere between 9.25% and 12.37%). In other words, Kevin Anthony guessed about right, and the computation in Example 7.2 is realistic.

The second column of Exhibit 7.3 presents the variability of forward market returns [fmr(t), see Equation (7.4)]. Note that only the first two lines are returns from the perspective of a U.S. investor; for the other currency pairs, we follow the usual conventions, so that the

Exhibit 7.3 Standard Deviations of Monthly Exchange Rate Changes and Forward Market Returns

Exchange Rate	Standard Deviation	
	Exchange Rate % Change	Forward Market Return
\$/€	11.17	11.25
\$/£	10.57	10.70
¥/\$	11.66	11.81
¥/€	11.34	11.42
£/€	9.25	9.35
¥/£	12.37	12.49

Notes: The table uses data from February 1976 to April 2010. The DEM replaces the euro before January 1999. The exchange rate % change is $s(t+1) = \frac{S(t+1) - S(t)}{S(t)} \times 100$ and the forward market return is $100 \times [s(t+1) - fp(t)]$ with $fp(t) = \frac{F(t) - S(t)}{S(t)}$. We annualize the monthly standard deviations by multiplying by the square root of 12, as is typical in financial markets. The data were obtained from Reuters and Global Insight.

investor is either yen- or pound-based. The variability of the forward market returns is of the same order of magnitude as the variability of the exchange rate changes themselves. Clearly, speculating in the foreign exchange market is not without risk of loss.

7.2 UNCOVERED INTEREST RATE PARITY AND THE UNBIASEDNESS HYPOTHESIS

Covered interest rate parity maintains that a domestic money market investment and a foreign money market investment have the same domestic currency return as long as the foreign exchange risk in the foreign money market investment is “covered” using a forward contract. Two related theories predict what may happen when exchange rate risk is, by contrast, not hedged. **Uncovered interest rate parity** maintains that the “uncovered” foreign money market investment, which has an uncertain return because of the uncertainty about the future value of the exchange rate, has the same *expected* return as the domestic money market investment. The **unbiasedness hypothesis** states that there is no systematic difference between the forward rate and the expected future spot rate and that, consequently, the expected forward market return is zero. In this section, we develop both of these hypotheses in more detail.

Uncovered Interest Rate Parity

If we take the expected value of the return to investing 1 dollar in the pound money market, as described in Equation (7.1), we find

$$E_t[r(t+1)] = \frac{1}{S(t)} \times [1 + i(\text{£})] \times E_t[S(t+1)]$$

Because the current spot rate, $S(t)$, and the interest rate, $i(\text{£})$, are in the time t information set, the expectation applies only to the future exchange rate.

Uncovered interest rate parity is the hypothesis that the expected return on the uncovered foreign investment equals the known return from investing 1 dollar in the dollar money market $[1 + i(\$)]$. If uncovered interest rate parity is true, there is no compensation to the uncovered investor for the uncertainty associated with the future spot rate, and expected returns

on investments in different money markets are equalized. Equivalently, the speculative return on borrowing 1 dollar and investing it in the pound money market, $\text{exr}(t+1)$ [see Equation (7.2)], is expected to be zero, given current information.

Let's go back to the portfolio manager Kevin Anthony. The interest rate on the pound is 12%, but the interest rate on the dollar is only 8%. Uncovered interest rate parity suggests that it would be naïve to think that pounds therefore constitute a great investment for Kevin. In fact, the high yield on pounds implies that the market anticipates the pound to depreciate by just enough that the expected dollar return to currency speculation in the pound market is also 8%. In particular,

$$\frac{1}{\$1.60/\pounds} \times [1 + 0.12] \times E_t[S(t+1)] = 1 + 0.08$$

or

$$E_t[S(t+1)] = \frac{1.08}{1.12} \times 1.60 = \$1.5429/\pounds$$

That is, the pound is expected to depreciate by 3.57%:

$$\left(\frac{\$1.5429/\pounds - \$1.60/\pounds}{\$1.60/\pounds} \right) = -0.0357$$

The Unbiasedness Hypothesis

When the forward rate equals the expected future spot rate, the forward rate is said to be an **unbiased predictor** of the future spot rate. This equality is summarized by the unbiasedness hypothesis:

$$F(t, \$/\pounds) = E_t[S(t+1, \$/\pounds)] \quad (7.5)$$

Covered interest rate parity and uncovered interest rate parity imply the unbiasedness hypothesis, which can be seen as follows (with S and F always referring to $\$/\pounds$ exchange rates):

$$E_t \left[\underbrace{\frac{S(t+1)}{S(t)}}_{\text{Uncovered Interest Rate Parity}} [1 + i(\pounds)] = [1 + i(\$)] = \underbrace{\frac{F(t)}{S(t)}}_{\text{Covered Interest Rate Parity}} [1 + i(\pounds)] \right] \quad (7.6)$$

By eliminating $S(t)$ and $[1 + i(\pounds)]$ from both sides of the exterior equations, we recover the unbiasedness hypothesis. To better understand the concept of an unbiased prediction, we must first understand the concept of a forecast error.

Forecast Errors

Whenever you predict something that is uncertain, such as the future spot rate, there will inevitably be a forecast error. A **forecast error** is the difference between the actual future spot exchange rate and its forecast. One way to measure the magnitude of forecast errors is to examine their standard deviation. We cannot just measure the average forecast error because very large errors in either direction would tend to cancel each other out, potentially resulting in a small average error. Because the standard deviation squares the errors, large errors result in a large standard deviation. In Exhibit 7.3, we showed that percentage changes in exchange rates and forward market returns are very variable. This large variability suggests that the forecast errors in predicting exchange rates, using either the current exchange rate or the forward rate as the forecast, are very variable. Forecasts from commercial firms that sell exchange rate forecasts also have large standard deviations, and no one forecasting firm seems to be very successful over time.

The Soros Saga

Of course, we do hear stories of speculators periodically making a fortune in the foreign exchange market. For example, the hedge funds operated by George Soros reportedly made \$2 billion in 1992, when Soros bet correctly that the British pound would weaken relative to the Deutsche mark. Soros subsequently became known as “the man who broke

the Bank of England.” What is less widely well known is that some years later, Soros lost over \$1 billion because he incorrectly bet that the euro would strengthen relative to the dollar. This and other difficulties eventually led Soros to change his strategy and make more conservative, safer investments.

Is it reasonable to expect exchange rate forecasts to be characterized with large variability? We think the answer is yes because exchange rates are the relative prices of currencies, and currencies are assets. Thus, exchange rates are asset prices, and we should expect exchange rates to behave very much like other asset prices, such as stock prices, which are also very difficult to predict. If exchange rates were easy to predict, lots of easy money would be made betting that one currency would strengthen relative to another.

Unbiased Predictors

An unbiased predictor implies that the expected forecast error is zero. In our setting, we forecast the future spot rate using the forward rate so that the forecast error is the difference between the two: $S(t+1) - F(t)$. The unbiasedness hypothesis states nothing about the magnitude of the forecast errors, which can be large or small and can vary over time. Instead, it has two important implications. First, given your current information, you should expect the forecast error to be zero. Second, on average, the forecast errors of an unbiased predictor may sometimes be negative and sometimes positive, but they are not systematically positive or negative, and they will average to zero.² If a forecast is biased, however, and you know what the bias is, you can improve your forecast by taking into account the bias. Currency speculators seek to exploit such biases.

The Unbiasedness Hypothesis and Market Efficiency

The unbiasedness hypothesis in Equation (7.5) is often identified with market efficiency. In efficient capital markets, investors cannot expect to earn profits over and above what the market supplies as compensation for bearing risk. An inefficient market is one in which profits from trading are not associated with bearing risks and are therefore considered extraordinary. The definition of **market efficiency** incorporates the hypotheses that people process information rationally and that they have common information on relevant variables that may help predict exchange rates. Together, these assumptions ensure that people have common expectations of the future.

To link the unbiasedness hypothesis more explicitly with market efficiency, recall the example of Mr. Buy and Ms. Sell. Mr. Buy's profit or loss from purchasing pounds forward, $S(t+1) - F(t)$, was equal but opposite in sign to Ms. Sell's profit or loss from selling pounds forward, $F(t) - S(t+1)$.

Notice that if the forward rate were a biased predictor of the future spot rate, and people had the same expectation of the future spot rate, one side of the forward contract, either Mr. Buy or Ms. Sell, would expect a profit on the contract, and the other party to the forward contract would expect a loss. Hence, the argument goes, because no one would willingly enter into a forward contract if they expected to lose money, forward rates must be unbiased

²The second implication follows from the first because of a famous statistical theorem called the Law of Iterated Expectations.

predictors of future spot rates if the market is efficient. That is, both Mr. Buy and Ms. Sell must both expect zero profits:

$$E_t[S(t+1, \$/\pounds) - F(t, \$/\pounds)] = 0 = E_t[F(t, \$/\pounds) - S(t+1, \$/\pounds)]$$

The unbiasedness hypothesis does run into a consistency problem when viewed from two different currency perspectives simultaneously. If it holds in dollars per pound, it must be violated when viewed from pounds per dollar. Appendix 7.1 analyzes this so-called Siegel paradox, demonstrating that it is not important in practice.

Uncovered interest rate parity and the unbiasedness hypothesis do take a narrow view of market efficiency, however. Because currency speculation involves risk taking, isn't it conceivable that there is a positive expected return to be made from speculating in the foreign exchange market? As long as the expected return is commensurate with the risk taken, earning an expected return would not be inconsistent with market efficiency.

7.3 RISK PREMIUMS IN THE FOREIGN EXCHANGE MARKET

You might be surprised by the fact that many rational people, like either Mr. Buy or Ms. Sell, are quite willing to enter contracts expecting a loss. Consider the purchase of fire insurance. Suppose you want to buy fire insurance for 1 year on your \$250,000 home. The insurance company charges you today and promises to pay you in the future if you suffer a certain type of loss—in this case, loss due to fire. Suppose everyone agrees that the probability of fire destroying your home is 0.1%. What insurance premium would you be willing to pay? If you are risk neutral, you would just be willing to pay the expected loss:

$$\$250,000 \times \frac{1}{100} \times 0.1 = \$250$$

However, if you confronted many people with this question, they would be willing to pay more than \$250 because they are risk averse. If they do, they willingly enter a contract with an expected loss because the expected value of the insurance (given the probability of a fire) is only \$250.

Similarly, going back to our earlier example, either Mr. Buy or Ms. Sell may be paying the other person a **risk premium** in order to avoid further harm from large exchange rate movements. For example, Ms. Sell may be selling pounds forward because she is the treasurer of a large multinational corporation (MNC) that is expecting future pound revenues. Remember that the forward rate is \$1.5429/£. Even if Ms. Sell expects the future spot rate to be higher than \$1.5429/£, she might still choose to hedge because there is a lot of uncertainty about the future value of the pound.

What Determines Risk Premiums?

The risk premium on an asset is the expected return on the asset in excess of the return on a risk-free asset. In this case, the excess return can be thought of as the uncovered foreign money market return, which we called $\text{exr}(t+1)$. Denoting the foreign exchange risk premium by rp , we have $rp(t) = E_t[\text{exr}(t+1)]$. Different assets can have different risks, and assets that are riskier must offer higher expected returns in order to induce investors to hold them. You may think that the riskiness of an investment in an asset is determined by the uncertainty associated with the asset's payoff. For example, the risk premium on currency speculation must be linked to the variability of exchange rate changes. After all, the conditional distribution of the future exchange rate will be wider the more variable such changes are. However, this is not the case. The reason is that investors care about the expected return

and risk of their whole portfolio of assets, not necessarily about the risk of an individual asset viewed in isolation.

Modern portfolio theory postulates that risk-averse investors like high expected returns on their portfolios, but they dislike a high variance in their portfolios. (That is, they don't like the value of their portfolios to go up and down very much.³) The question then becomes: How does the return on an individual asset contribute to the variance of the kinds of portfolios investors are likely to hold? It turns out that if there are many assets in the portfolio, part of the variance of an asset's return does not contribute to the portfolio's variance. This leads to an important decomposition of the uncertainty of the return on any asset.

Systematic and Unsystematic Risk

The uncertainty of a return can always be decomposed into a part that is *systematic* and a part that is *unsystematic*, which is also called *idiosyncratic*. That is,

$$\text{Individual asset return uncertainty} = \text{Systematic risk} + \text{Idiosyncratic uncertainty}$$

Systematic risk is the risk associated with an asset's return arising from the covariance of the return with the return on a large, well-diversified portfolio. The **covariance** of two random variables describes how the two variables move together, or *covary*, with each other. Often, we describe how things covary with each other in terms of **correlation** coefficients that are bounded between -1 and $+1$. If the returns on two assets are perfectly correlated (that is, they always perfectly move in the same direction), their correlation coefficient is 1 . By contrast, if the assets are not at all correlated (neither moves at all in relation to the other), their correlation coefficient is 0 . If the coefficient is -1 , the two asset returns always move in opposite directions. The correlation coefficient is the covariance of the two variables divided by the product of their standard deviations.

The large, well-diversified portfolio that investors should hold according to finance theory is called the **market portfolio**.⁴ The market portfolio is the value-weighted collection of all available financial assets in the market as a whole.

How does this decomposition relate to risk premiums? If the return on the asset contains only idiosyncratic uncertainty, there will be no increase in the expected return on the asset due to the uncertainty of the return. It will not command a risk premium! The asset will be priced to yield an expected return equal to the return on risk-free assets. An asset has only idiosyncratic uncertainty if its return does not covary with the returns on other assets.

These statements follow from a fundamental insight of portfolio theory: Idiosyncratic uncertainty can be diversified away. Even though investors do not like the uncertainty of their total portfolio and demand risk premiums on assets that contribute to the variance of the portfolio, assets whose returns contain only idiosyncratic uncertainty do not contribute to the variance of the portfolio and, consequently, do not command any risk premium. Because idiosyncratic uncertainty is diversifiable in large portfolios, it is also called *diversifiable uncertainty*, or *diversifiable risk*. Because systematic risk measures how much an asset's return co-moves with the market, it cannot be diversified away, and the risk involved commands a risk premium. For example, the variance of an individual stock return is partly driven by macroeconomic events such as the business cycle and interest rates that affect every stock. Such risks are systematic. The variance of the stock return is also partially driven by **idiosyncratic risks** that affect only that particular stock, such as the quality of the firm's management.

³We have previously discussed the variance of a random variable and indicated that it is a measure of the dispersion of the probability distribution. Graphically, the square root of the variance (the standard deviation) is associated with the width of a bell-shaped curve.

⁴Appendix 7.2 provides a review of portfolio theory and related statistical concepts, such as covariance, correlations, and betas, to allow you to examine the arguments implying that covariances among returns are the main sources of portfolio variance.

The CAPM

The theories we have been discussing are the foundation of the **capital asset pricing model (CAPM)**. William F. Sharpe was awarded the Nobel Prize in Economics in 1990 for its development. The CAPM holds that it is the covariance of an asset's return with the return on the market portfolio that determines the asset's systematic risk and hence its risk premium. The model also provides an easy-to-implement procedure to put an actual number on the risk premium, which we describe in detail in Chapter 13.

According to the CAPM, the systematic risk of an individual asset is fully described by its **beta** with respect to the market portfolio. The formula for the beta is simple:

$$\text{Beta on asset } i = \frac{\text{Covariance (Asset return } i, \text{ Market portfolio return)}}{\text{Variance (Market portfolio return)}}$$

Higher betas indicate higher systematic risk, and the CAPM postulates that

$$\text{Risk premium on asset } i = (\text{Beta on asset } i) \times (\text{Risk premium on market portfolio})$$

What is the intuition for the prediction about expected returns of the CAPM? Think of the return on the risk-free asset as the compensation provided to an investor for the time value of money that is required by the investor because the investor sacrifices the use of the money for a certain period. The investor requires compensation in excess of the risk-free rate (that is, a risk premium) if the beta of the asset is positive, as are the betas of most equity investments. Assets with positive betas contribute to the variance of the market portfolio, and the larger the beta, the riskier the asset and the higher its expected return must be. Notice that if an asset has a negative beta because the return on the asset is negatively correlated with the return on the market portfolio, the expected return on the asset is less than the risk-free rate. Investing in an asset that covaries negatively with the return on the market portfolio provides an investor with portfolio insurance. When the rest of the investor's portfolio is doing poorly, the asset with the negative covariance generally pays high returns, and when the rest of the investor's portfolio pays high returns, the asset with the negative covariance generally pays relatively low returns. Investing in this asset thus dampens the volatility of the return on the total portfolio. Risk-averse investors are willing to "pay" for this reduction in the volatility of their overall portfolio by accepting an expected return that is less than the risk-free interest rate.

Applying the CAPM to Forward Market Returns

Because a forward contract is an asset, there is potential for a risk premium. How will this bias the forward rate as a predictor of the future spot rate? Taking a position in a forward contract involves no investment of funds at the point in time when the contract is set, and it is not necessary to compensate the investor for the time value of money. But the dollar profits and losses on the forward contract can still covary systematically with the dollar return on the market portfolio. Hence, if the profitability of Mr. Buy's purchase of foreign currency at the forward exchange rate covaries positively with the dollar return on the market portfolio, Mr. Buy will view the forward contract as risky and will demand an expected profit. As noted previously, though, Ms. Sell's profits and losses on the forward contract are the opposite of Mr. Buy's profits and losses. Hence, if Mr. Buy's dollar profit is positively correlated with the dollar return on the market portfolio, the covariance of the dollar profit on Ms. Sell's side of the forward contract is negatively correlated with the dollar return on the market portfolio. In this case, when Ms. Sell enters into the contract, she obtains an asset that reduces the variability of her overall portfolio. She consequently willingly holds this contract at an expected loss. Again, this is like portfolio insurance. From Ms. Sell's perspective, the expected loss is balanced by the fact that the forward contract performs well when the rest of her portfolio does poorly. Consequently, there can be a risk premium that causes the forward rate to be a biased predictor of the future spot rate. According to the CAPM, such a risk premium should depend on the beta of the (excess) return to currency speculation.

Formal Derivation of CAPM Risk Premiums (Advanced)

The CAPM in Symbols

Let the dollar return for a 1-year holding period for an arbitrary asset j be $R_j(t+1)$, and let the risk-free return be $[1 + i(t, \$)]$. The CAPM predicts that the risk premium on an asset is equal to the beta of the asset multiplied by the amount by which the expected return on the market portfolio, $R_M(t+1)$, exceeds the return on the risk-free asset:

$$E_t\{R_j(t+1) - [1 + i(t, \$)]\} = \beta_j E_t\{R_M(t+1) - [1 + i(t, \$)]\} \quad (7.7)$$

The beta of the j th asset is the covariance of the return on asset j with the return on the market portfolio, σ_{jM} , divided by the variance of the return on the market portfolio, σ_{MM} :

$$\beta_j = \frac{\sigma_{jM}}{\sigma_{MM}}.$$

Here, the variance (covariance) is a conditional variance (covariance) because it is based on the information at time t .

The CAPM and Forward Market Returns

Let's derive the implications of the CAPM for the risk premium on an unhedged investment of dollars in the British pound money market. The uncovered excess return was defined in Equation (7.2), and we review it here for convenience:

$$\text{exr}(t+1) = \frac{S(t+1, \$/\pounds)[1 + i(t, \pounds)]}{S(t, \$/\pounds)} - [1 + i(t, \$)] = R_{\pounds}(t+1) - [1 + i(t, \$)]$$

From Equation (7.7), the CAPM gives the expected excess return on this uncertain dollar investment:

$$E_t[\text{exr}(t+1)] = \beta_u E_t\{R_M(t+1) - [1 + i(t, \$)]\} \quad (7.8)$$

The beta on the uncovered pound investment is

$$\beta_u = \frac{\text{COV}_t[R_{\pounds}(t+1), R_M(t+1)]}{\text{VAR}_t[R_M(t+1)]}$$

where COV_t and VAR_t are shorthand for conditional covariance and variance, respectively, and the interest rates do not enter the expression because they are in the time t information set.

The forward market return also satisfies a CAPM relationship:

$$E_t[\text{fmr}(t+1)] = \beta_F E_t\{R_M(t+1) - [1 + i(t, \$)]\} \quad (7.9)$$

Here, β_F is the beta on the forward contract to buy foreign currency in the forward market and sell it subsequently in the future spot market. Recall from Section 7.1 that $\text{fmr}(t+1) = \frac{\text{exr}(t+1)}{1 + i(t, \pounds)}$. Therefore, $\beta_F = \frac{\beta_u}{1 + i(t, \pounds)}$. In other words, the expected returns on forward market contracts and money market investments are proportional because they have the same fundamental risk exposure but invest a different number of units.

Equations (7.8) and (7.9) indicate that forward rates will be biased predictors of future spot rates if there is systematic risk associated with the profits on a forward contract. In the case of the dollar/pound example, if the dollar weakens relative to the pound when the dollar payoff on the market portfolio is high, the risk premium would be positive, and the forward rate would be below the expected future spot rate. You would expect to profit by buying pounds forward, and you would expect to suffer a loss by selling pounds forward. If, on the other hand, the dollar strengthens relative to the pound when the dollar return on the market

portfolio is high, the beta on the forward contract would be negative. Thus, the forward rate would be above the expected future spot rate, and there would be an expected loss from buying pounds forward and an expected gain from selling pounds forward.

7.4 UNCOVERED INTEREST RATE PARITY AND THE UNBIASEDNESS HYPOTHESIS IN PRACTICE

Taking a stand on whether uncovered interest rate parity and the unbiasedness hypothesis actually hold is important when international financial managers make decisions. This section reviews situations in which this issue arises.

Situations Where Premiums Matter

International Portfolio Management

When a European portfolio manager buys Japanese equities, he hopes the Japanese equity market will perform well, but he is also exposed to foreign exchange risk in the yen–euro market. As we discuss in detail in Chapter 13, the return on a foreign bond and/or equity can be decomposed into two components: the (local) return on the foreign asset and the currency return. Global money managers may decide to speculate on a currency, or they may decide to hedge the currency risk. This decision is greatly affected by whether they believe in the validity of uncovered interest rate parity and the unbiasedness hypothesis.

The Cost of Hedging

Multinational corporations often hedge their transaction foreign exchange risk using forward contracts. Clearly they may be willing to pay a premium to insure against this risk. The following *Point–Counterpoint* makes a link between the unbiasedness hypothesis and a practical hedging situation. In a nutshell, when unbiasedness holds, multinationals effectively do not pay premiums to hedge their transaction foreign exchange risk. Of course, as we argued in Section 7.3, the existence of a premium is not necessarily inconsistent with market efficiency and may be fair compensation for risk insurance. Note also that an MNC may benefit from such premiums. For example, if the long position in a particular currency commands a premium, an MNC that hedges a short position will earn the risk premium.

Exchange Rate Forecasting

Forecasting exchange rates is difficult, but it remains an activity that attracts many resources and much brainpower in the real world. If the unbiasedness hypothesis holds, the best forecast of the future exchange rate can be read from a table in your daily *Financial Times* or *Wall Street Journal* because the answer lies in the forward rate. Chapter 10 examines the success of different forecasting models relative to the forward rate.

Exchange Rate Determination

Chapter 10 discusses some popular exchange rate determination theories. It turns out that many of the well-known theories linking exchange rate values to fundamentals such as trade balances, money supplies, and so forth, assume that uncovered interest rate parity holds. But if it does not hold, the validity of these theories is immediately in doubt. On the other hand, the empirical evidence that we present in Section 7.5 has motivated some macroeconomists to supplement macro-models with time-varying foreign exchange risk premiums.

POINT-COUNTERPOINT

The Cost of Hedging⁵

Ante and Freedy's Uncle Fred is holding forth during dinner at the annual Handel family gathering at his estate in Chappaqua, New York. Uncle Fred is in the export–import business, is very well traveled, and loves recounting his on-the-road war stories. After a hilarious account of how a Dutch business associate recommended checking out the Walleijes (the red light district) in Amsterdam as the high point of Dutch architecture, he suddenly turns to Ante and Freedy: “Hey, how’s that international finance class going? I hope well, because I’ve got a question for you from my business. Suppose I owe 10 million Swedish kronor, payable in 1 month. My company has the cash to buy kronor now, or it could wait until later. I figure we should put the money wherever in the world it would earn the highest interest rate, but my treasurer, an MBA hotshot, tells me that high interest rates are irrelevant because if the krona interest rate is higher than the dollar interest rate, the krona is expected to fall in value relative to the dollar. When I ask her what I should do, she says that it doesn’t matter. ‘Flip a coin,’ she says. Is this why I’m paying her such a high salary? Anyway, young fellows, what do you think?”

As usual, Ante is quickest to respond: “You’re absolutely right, Uncle Fred, you should fire that MBA. I am convinced that you will earn a higher return if you put your cash balances in the currency that has the highest interest rate. That way, you will lower the effective dollar cost of your foreign payables.”

Freedy shakes his head. “Have you been sleeping in class, Ante? Remember the theory of uncovered interest rate parity? The MBA is right. On average, dollar returns will be equalized in different countries. If Uncle Fred puts his money in kronor when the interest rate is high, the krona will likely depreciate, wiping out the interest rate gain. Maybe he could make it easier on himself and just buy the kronor in the forward market.”

“Hmm, this is a useful argument. Let’s have our grappa in the living room. Maybe that will bring your thoughts together,” sighs Uncle Fred. As they walk toward the comfortable, Italian-designed sofas, Suttle Trooth joins them from the kitchen.

“Hey guys, I overheard your conversation, and are you ever confused,” says Suttle. “Let me explain to Uncle Fred what is going on. I brought some paper and a pencil because I want to write down a few things.”

“Consider what Uncle Fred is saying,” continues Suttle. “Suppose he keeps his money in dollars. Then, Uncle Fred incurs currency risk because he will have to convert the dollars into kronor 1 month from now at the exchange rate of $S(t+1, \$/\text{SEK})$. The dollar cost in 1 month of the krona payable will be

$$\text{SEK}10 \text{ million} \times S(t+1, \$/\text{SEK})$$

If he converts his dollars now, he will not face any currency risk because he will know exactly how many kronor to convert so that they grow to SEK10 million in 1 month. That amount will be the present value of the SEK10 million, or

$$\text{SEK}10 \text{ million} \times \frac{1}{1 + i(\text{SEK})}$$

The current dollar cost of this amount of kronor is

$$\text{SEK}10 \text{ million} \times \frac{1}{1 + i(\text{SEK})} \times S(t, \$/\text{SEK})$$

⁵This *Point-Counterpoint* is motivated by the discussion in Kenneth Froot and Richard Thaler (1990).

Because the first cost is dollars in 1 month and the second cost is dollars today, to compare the alternative strategies, we have to take both costs to the same point in time. Taking the future value in dollars of the second strategy gives

$$\text{SEK10 million} \times \frac{1}{1 + i(\text{SEK})} \times S(t, \$/\text{SEK}) \times [1 + i(\$)]$$

At this point, Freedy interjects, “Hey, those terms involving interest rates and the spot rate are equal to the forward rate, right?”

“Very good, Freedy, you’ve got it,” replies Suttle. “The strategy of converting into kronor now is equivalent to a strategy of buying kronor in the forward market. Therefore, we can compare the performance of Uncle Fred’s possible strategies by comparing the future exchange rate with the forward rate. Suppose dollar interest rates are higher than krona interest rates, in which case the krona trades at a forward premium. Then, Uncle Fred’s proposal would have him not hedge, and he would keep his money in dollars. That strategy works great *if* the future USD/SEK exchange rate turns out to be lower than the forward rate. If it does, Uncle Fred’s *ex post* costs will be relatively low.”

“Very interesting, but all these equations do not appear to answer my question, now do they?” grumbled Uncle Fred.

“Hold on. I am not done yet,” says Suttle. “Let’s think about what you’d lose by hedging. We can call this the *cost of hedging*, if you wish. *Ex post*, the cost of having hedged can be either positive or negative because it will equal

$$F(t, \$/\text{SEK}) - S(t+1, \$/\text{SEK})$$

If the forward rate is higher than the future spot rate, you would indeed have been better off not to hedge and to have just taken the currency risk. Of course, you cannot necessarily know when this will occur, and there will certainly be instances in which the future spot rate ends up higher than the forward rate (when the SEK appreciates more than the forward rate indicates), in which case your *ex post* cost of hedging will be negative because you have higher costs by having not hedged. Now, what the MBA is trying to tell you is that the expected value of the cost of hedging is zero in an efficient market with no risk premium:

$$E[F(t, \$/\text{SEK}) - S(t+1, \$/\text{SEK})] = 0$$

This relationship is also known as the unbiasedness hypothesis. Equivalently, whether interest rates are higher or lower abroad does not matter because currency changes, on average, correct for this. If the unbiasedness hypothesis is correct, it won’t matter whether you hedge or do not hedge your exposure. Also, Uncle Fred, your strategy won’t make money on average because sometimes you will hedge and sometimes you will not, but the expected difference between the two is zero. So the expectation of the difference in the cost of the two strategies can be viewed as the expected cost of hedging, and it is zero—if unbiasedness holds.”

Ante excitedly interjects, “But who says the market is efficient? These equations are derived by some ivory tower academics. Why should we expect them to characterize actual markets where real people have to trade?”

“Well, there is something to that point, I must admit,” answers Suttle. “Some econometric tests have rejected the unbiasedness hypothesis, and the estimates actually indicate that Uncle Fred’s high-yield strategy may work. But that need not mean the market is inefficient. If Uncle Fred does not hedge, he is exposed to currency risk. In other asset markets, such as equities, investors are compensated for taking on risk by receiving a higher expected return than the risk-free rate. We call this higher expected return a *risk premium*. There are probably risk premiums in the currency markets, too. If indeed there is a risk premium, there is an expected cost or an expected return to hedging. Suppose that a relatively high interest rate is providing compensation for both expected currency depreciation but also for risk. Uncle Fred’s unhedged

strategy is then associated with currency exposure when such exposure is very risky. To make this more concrete, suppose the dollar interest rate is higher than the krona interest rate. Uncle Fred won't hedge because he thinks $E[F(t) - S(t+1)] > 0$. There is a positive cost to hedging. But is that wise? Uncle Fred is not in the foreign exchange investment business, exchange rates are quite volatile, and not hedging may really hurt the bottom line, if the currency moves against him. When you hedge, you buy security! Don't you agree, Uncle?" asks Suttle, turning to see Uncle Fred comfortably snoring on the Italian sofa.

7.5 EMPIRICAL EVIDENCE ON THE UNBIASEDNESS HYPOTHESIS

In this section, we derive statistical tests of whether forward rates have historically been unbiased predictors of future spot rates and apply them to exchange rate data. The discussion uses basic statistics reviewed in Chapter 3 and **regression analysis**. (Appendix 7.3 provides a primer on regression tests.)

The Quest for a Test

A proper econometric test of the unbiasedness hypothesis transforms Equation (7.5) by dividing by $S(t, \$/\pounds)$ on both sides and by subtracting 1—with 1 written as $S(t, \$/\pounds)/S(t, \$/\pounds)$ —from both sides of the equation.⁶ This is possible because the spot exchange rate at time t , $S(t, \$/\pounds)$, is in the investors' information set.

$$\begin{aligned} fp(t, \$/\pounds) &\equiv \frac{F(t, \$/\pounds) - S(t, \$/\pounds)}{S(t, \$/\pounds)} \\ &= \frac{E_t[S(t+30, \$/\pounds) - S(t, \$/\pounds)]}{S(t, \$/\pounds)} \equiv E_t[s(t+30, \$/\pounds)] \end{aligned} \quad (7.10)$$

In Equation (7.10), we use a 30-day (1-month) forward contract, as in the empirical test reported in the next section. The left-hand side of Equation (7.10) is recognized as the 30-day forward premium (fp) or discount on the pound. The right-hand side of Equation (7.10) is the expected rate of appreciation or depreciation of the pound relative to the dollar (s). Equation (7.10) states that the unbiasedness hypothesis requires the forward premium or discount on the pound to be equal to the market participants' expectations about the rate of appreciation or depreciation of the pound relative to the dollar over the course of the next 30 days. If the hypothesis holds, the expected return to currency speculation will be exactly zero.

Incorporating Rational Expectations into the Test

The most difficult problem in testing the unbiasedness hypothesis is that it contains a variable that cannot be observed by a statistician: the conditional expectation of the rate of appreciation of the pound relative to the dollar. This conditional expectation is formed by market participants on the basis of their information set. Hence, in order to test the unbiasedness hypothesis, a statistician must specify how investors and speculators form their expectations. Typically, when statisticians are confronted with an unobservable variable, they make an auxiliary assumption to develop a test of the underlying hypothesis.

As in most other areas of financial economics, the most popular auxiliary assumption is that investors have **rational expectations**. If investors have rational expectations, they do not

⁶Because spot rates and forward rates move together over time in a very persistent fashion, a test in levels of the variables would almost always fail to reject the unbiasedness hypothesis, even when the hypothesis was false (see Engel, 1996).

make systematic mistakes, and their forecasts are not systematically biased. When investors have rational expectations, we can decompose the realized (observed) rate of appreciation into its conditional expectation plus an error term that does not depend on time t information:

$$s(t+30, \$/\pounds) = E_t[s(t+30, \$/\pounds)] + \varepsilon(t+30) \quad (7.11)$$

Realized appreciation = Expected appreciation + Forecast error

The error term can be viewed as news that moved the exchange rate, but the news, by definition, was unanticipated by rational market participants at time t .

Rational expectations imply that both the conditional mean, $E_t[\varepsilon(t+30)]$, and unconditional mean, $E[\varepsilon(t+30)]$, of the error term, $\varepsilon(t+30)$, in Equation (7.11), are zero. Because it reflects unanticipated news, $\varepsilon(t+30)$ should not be correlated with anything in the information set. Substituting the unbiasedness hypothesis of Equation (7.10) into Equation (7.11), we obtain

$$s(t+30, \$/\pounds) = fp(t, \$/\pounds) + \varepsilon(t+30) \quad (7.12)$$

In Equation (7.12), one observable variable, the realized rate of appreciation, equals another observable variable, the forward premium, plus an unobservable error term whose conditional mean is zero. This equation can be used for two tests of the unbiasedness hypothesis.

A Test Using the Sample Means

Because the average or mean forecast error in Equation (7.12) should be zero, we can easily test the weakest implication of the unbiasedness hypothesis: The unconditional mean of the realized rate of appreciation should equal the unconditional mean of the forward premium.⁷ The equality of these means or averages constitutes the **null hypothesis** (the hypothesis that is assumed to be true and is tested using data and a test statistic). Intuitively, to test the hypothesis, we compare the two sample means and check whether the difference between them is small or large in a statistical sense.

Data on Rates of Appreciation and Forward Premiums

The equality of the mean rate of appreciation and the mean forward premium is examined in Exhibit 7.4, which reports the results for all possible exchange rates between the dollar, the euro (the Deutsche mark before 1999), the British pound, and the Japanese yen. The data are expressed in annualized percentage terms. Consequently, the value of -2.82 for the mean rate of change of the dollar relative to the yen indicates that during the sample period, the dollar weakened relative to the yen at an average annual rate of 2.82%. The sample means of the realized rates of appreciation range from -3.70% for the yen value of the pound to 2.81% for the pound value of the euro. We can conclude that the mean of a time series is significantly different from zero at the 95% confidence level if the sample mean is more than 1.96 standard errors from zero. Said differently, we are then 95% sure that the true mean is not zero. The standard error of the sample mean depends on the volatility of the time series and the number of observations.⁸ In all cases, the volatilities of the rates of appreciation are large. The large volatility of the realized rate of appreciation inflates the standard errors associated with the mean rate of appreciation, making it difficult to precisely estimate the mean. Thus, not a single mean rate of depreciation is sufficiently large relative to its standard error that we can be more than 90% confident that it is significantly different from zero.

⁷The sample mean of a time series x_t using T observations is $\frac{1}{T} \sum_{t=1}^T x_t$.

⁸The usual standard error of the sample mean for a time series is σ/\sqrt{T} , where $\sigma^2 = \sum_{t=1}^T (x_t - \hat{\mu})^2/T$ denotes the sample variance of the series, and $\hat{\mu}$ denotes the sample mean of the series. For this to be the correct standard error, the time series must be serially uncorrelated, that is, the observation at time t must not be correlated with the observation at time $t+1$. The standard errors reported here are slightly different because they are calculated using the methods of Hansen (1982) and accommodate both serial correlation and conditional heteroskedasticity (see Chapter 2).

Exhibit 7.4 Means of Monthly Rates of Appreciation, Forward Premiums, and the Differences Between the Two

Exchange Rate	Mean		
	Rate of Appreciation (S.E.) Conf.	Forward Premium (S.E.) Conf.	Difference (S.E.) Conf.
\$/€	2.04	1.45	0.58
	(1.98)	(0.25)	(2.02)
	0.70	1.00	0.23
\$/£	-0.41	-2.24	1.62
	(1.93)	(0.23)	(1.99)
	0.17	1.00	0.59
¥/\$	-2.82	-3.31	0.48
	(2.05)	(0.23)	(2.13)
	0.83	1.00	0.18
¥/€	-1.44	-1.86	0.42
	(2.03)	(0.18)	(2.07)
	0.52	1.00	0.16
£/€	2.81	3.50	-0.69
	(1.64)	(0.25)	(1.67)
	0.91	1.00	0.32
¥/£	-3.70	-5.34	1.65
	(2.29)	(0.19)	(2.34)
	0.89	1.00	0.52

Notes: The table uses data from February 1976 to April 2010. Before 1999, the DEM replaces the euro. The monthly data are expressed as annualized percentage rates. The standard error (s.e.) measures the uncertainty we have about the accuracy of our estimate of the sample average. If we had an infinite amount of data, the standard error would be zero. As a technical note, the standard errors allow for conditional heteroskedasticity and two lagged autocorrelations in the errors. The confidence level (Conf.) of the test that the mean is zero is below the standard error. A confidence level of 0.90 indicates that we can be 90% sure that the null hypothesis of a zero mean is false.

The sample means of the forward premiums range from -5.34% for the yen value of the pound to 3.50% for the pound value of the euro. Because the volatilities of the forward premiums are much smaller than those of the rates of appreciation, all the sample means of the forward premiums are large relative to their respective standard errors. Hence, we can be quite confident that all the unconditional means of the forward premiums are not zero. For example, the pound appears robustly at a forward discount relative to all other currencies.

The Test

The third column of Exhibit 7.4 tests the hypotheses that the means of the 1-month forward premiums are equal to the means of the 1-month rates of appreciation on a currency-by-currency basis. The third column is labeled “Difference” to indicate that it represents the (*ex post*) rate of appreciation minus the (*ex ante*) forward premium. If the null hypothesis is true, the mean of the difference should be zero. In no case is there sufficient evidence to reject the null hypothesis with 90% confidence. The largest confidence level is only 0.59 for the dollar value of the pound. Of course, here again, the volatilities of the realized rates of appreciation make it difficult to precisely estimate the differences of the means.⁹

⁹Because of triangular arbitrage, only three of the six statistical tests we conducted provide independent information. When we do a joint test for the difference between the mean rate of appreciation of the euro relative to the three other currencies and the three corresponding average forward premiums, we also fail to reject that the differences are jointly zero.

In sum, there is essentially no evidence to suggest that the unconditional means of the forward premiums differ from the unconditional means of the rates of appreciation. Because the difference between $s(t+1)$ and $fp(t)$ is the forward market return, our test results imply that, on average, forward market returns are zero.

High-Interest-Rate Currencies Depreciate

The zero unconditional means of the differences between the rates of appreciation and the forward premiums are also consistent with an important fact of international finance: Countries with high nominal interest rates have currencies that tend to depreciate in value over time relative to the currencies of countries with low nominal interest rates. From our discussion of interest rate parity in Chapter 6, you know that the forward premium on a foreign currency is equal to the interest differential between the domestic currency and the foreign currency. Hence, failure to reject the unbiasedness hypothesis with the test of unconditional means supports the proposed fact quite strongly. For example, the average forward discount on the euro in terms of the yen is 1.86%, which implies that the euro (formerly DEM) interest rates were on average 1.86% higher than JPY interest rates. Exhibit 7.4 demonstrates that these higher euro interest rates were providing compensation for the average depreciation of the euro relative to the yen, which was 1.44%, not much smaller than 1.86%.

One interesting aspect of the differences reported in Exhibit 7.4 is that with the exception of the dollar/euro pair, the high-interest-rate currencies do appear to depreciate less than the forward discount indicates. In other words, forward market returns from long positions in weak currencies are, typically, on average positive. Lustig et al. (2009) and Jylhä et al. (2010) have argued that these positive returns for weaker currencies reflect risk premiums, either because these currencies are more exposed to global risk factors or because the inflation environment in these countries is riskier. Yet, Exhibit 7.4 suggests that the statistical evidence for these premiums remains weak.

In assessing the validity of the unbiasedness hypothesis, it is important to remember that this first test is a very weak implication because it considers only the overall average performance of the theory. We can also derive tests that examine the implications of the theory at different points in time. Such an approach is important because it corresponds to what someone would do in an active international portfolio management situation.

Regression Tests of the Unbiasedness of Forward Rates

A straightforward way to use additional information to test the unbiasedness hypothesis is to use regression analysis. Suppose we write Equation (7.12) in the form of a regression, as in the following equation:

$$s(t+30) = a + bfp(t) + \varepsilon(t+30) \quad (7.13)$$

Here, a is the intercept, and b is the slope coefficient of the regression. The unbiasedness hypothesis implies that $a = 0$ and $b = 1$ because with these substitutions, Equation (7.13) reduces to Equation (7.12).

The regression tests of the unbiasedness hypothesis are presented in Exhibit 7.5, which presents the estimated parameters and their standard errors for regressions using the same six exchange rates as in Exhibit 7.4. The standard errors are presented in parentheses below the estimated coefficients. The confidence levels of the tests that $a = 0$ and that $b = 1$ are presented below the standard errors. Values of the confidence level that are above 0.90 indicate that we can reject the null hypothesis with 90% confidence.

Exhibit 7.5 Regression Tests of the Unbiasedness Hypothesis
 $s(t+30) = a + b fp(t) + \varepsilon(t+30)$

Currency	Coefficients on Regressors		R^2
	Const. (S.E.)	Forward Premium (S.E.)	
	Conf. ($a = 0$)	Conf. ($b = 1$)	
\$/€	3.26	-0.84	0.004
	(2.31)	(0.81)	
	0.84	0.98	
\$/£	-3.84	-1.68	0.016
	(2.24)	(0.82)	
	0.91	1.00	
¥/\$	-10.03	-2.18	0.023
	(2.67)	(0.64)	
	1.00	1.00	
¥/€	-4.46	-1.62	0.008
	(2.30)	(0.87)	
	0.95	1.00	
£/€	4.70	-0.54	0.003
	(2.56)	(0.65)	
	0.93	0.98	
¥/£	-17.17	-2.52	0.020
	(5.34)	(0.84)	
	1.00	1.00	

Notes: The table uses data from February 1976 to April 2010. The euro replaces the Deutsche mark (DEM) from 1999 onward. Data on rates of appreciation and the forward premiums are annualized. The parameter estimates are obtained using ordinary least-squares regression for each equation. The standard error (s.e.) is in parentheses below the estimate. The confidence level (Conf.) of the test is below the standard error. The tests are that the constant is 0 and that the slope coefficient is 1. The last column reports the R^2 : how much of the variation in $s(t+30)$ is explained by the variation in $fp(t)$. The standard errors correct for heteroskedasticity and allow for serial correlation (2 lags) in the error terms.

Notice that all six of the estimated values of b are significantly different from unity. Perhaps more surprisingly, notice that all the estimated slope coefficients are negative. The estimated values of b range from -2.52 for the yen value of the pound to -0.54 for the pound value of the euro. Consequently, the regressions suggest the existence of a **forward rate bias**; the forward rate does not equal the expected future spot rate. The regression evidence thus qualifies the use of the unbiasedness hypothesis. Treasurers in MNCs and global portfolio managers must realize that there is a potential cost to hedging foreign currency risk because the forward rate is not necessarily the best forecast of the future exchange rate.

Because negative values of b are found in the cross-rate regressions as well, the explanation of this phenomenon for the dollar exchange rates should not be sought in a story about common movements of the dollar relative to other currencies, nor could it be due strictly to U.S. policy. Apparently, the explanation must encompass the behavior of all major foreign exchange markets.

Notice also that the explanatory power of the regressions, which is measured by the R^2 values, is quite low. The largest R^2 is 2.3%. The appropriate way to interpret this finding is that there is some ability of the forward premium to predict the rate of appreciation,

but the unanticipated component in the rate of appreciation is large relative to its predictable component.

Interpreting the Forward Bias

The unbiasedness regression generates a forecast for the future changes in exchange rates and hence also for the forward market return

$$E_t[s(t+1)] = \hat{a} + \hat{b}fp(t) \quad \text{or} \quad E_t[s(t+1) - fp(t)] = \hat{a} + (\hat{b} - 1)fp(t) \quad (7.14)$$

Note that $s(t+1) - fp(t)$ is nothing but the forward market return, the return to a long forward position in the foreign currency.

People familiar with the results of the unbiasedness regressions just presented often argue that the negative slope coefficients imply that currencies trading at a forward discount will strengthen, in contrast to the prediction of the unbiasedness hypothesis, which implies that discount currencies are going to weaken. Unfortunately, this interpretation of the regression is wrong because it ignores the value of the constant term in the regression.

Exhibit 7.6 shows the importance of the constant in the regression, using the yen/dollar equation as an example. We consider a forward discount on the dollar of 3.31%, the sample average (see Exhibit 7.4), implying that Japanese yen interest rates were on average approximately 3.31% less than U.S. dollar interest rates. On the first line of Exhibit 7.6, we repeat the prediction of the theory: If the dollar is at a 3.31% discount, it should be expected to depreciate by 3.31%. If we were to use the regression and ignore the constant as in the computation on the second line, the prediction is a 7.22% appreciation of the dollar, so that the dollar indeed gives a higher yield and is expected to appreciate substantially.

However, the correct interpretation is on the third line of Exhibit 7.6, which uses the regression with the estimated coefficients as in Equation (7.6) to determine an estimate of expected dollar depreciation or appreciation. The dollar is now expected to weaken, but only by 2.82%. This is the average depreciation of the dollar over the sample period (see Exhibit 7.4), and most importantly, it is lower than the depreciation the forward discount suggests. However, the regression still implies that a speculator should buy dollars forward if he believes the prediction of the regression will be borne out. That is,

$$\begin{aligned} E_t[\text{fmr}(t+1)] &= E_t[s(t+1) - fp(t)] \\ (\text{Expected forward market return}) &= -2.82\% - (-3.31\%) \\ &= 0.49\% \end{aligned}$$

Exhibit 7.6 Interpreting the Unbiasedness Regression

	$fp(t)$	a	b	$E_t[s(t+1)]$	$E_t[\text{fmr}(t)]$
Uncovered Interest Rate Parity	-3.31%	0	1	-3.31%	0%
Naïve Interpretation	-3.31%	0	-2.18	7.22%	10.53%
Actual Interpretation	-3.31%	-10.03	-2.18	2.82%	0.49%
(large discount)	-5.00%	-10.03	-2.18	-0.87%	5.87%

Notes: The four different lines compare expected exchange rate appreciation using information in the forward premium and three different assumptions. The first line assumes uncovered interest rate parity holds. The second line uses the regression reported in Exhibit 7.5 for ¥/\$ but sets the constant equal to 0. The third line uses the actual regression results. In the fourth line, we consider a larger forward discount. All the percentages are annualized.

The expected forward market return from buying dollars forward is positive! When the forward discount is unusually large, there can be an expected dollar appreciation, and the expected return from going long dollars increases substantially. The last line in Exhibit 7.6 demonstrates this for a forward discount of 5%.

7.6 ALTERNATIVE INTERPRETATIONS OF THE TEST RESULTS

In this section, we examine three possible explanations of the results from the preceding section: market inefficiency, the presence of a foreign exchange risk premium, and peso problems.

Market Inefficiency

The evidence against the unbiasedness hypothesis suggests that interest rate differentials may contain information about future exchange rates that can be profitably exploited. Both academic analysts and foreign exchange professionals have explored models that link future exchange rate changes to interest rate differentials and other easily available information (such as past exchange rates) to predict future exchange rates (see, for example, Villanueva, 2007).

Exploiting the Forward Bias and Carry Trades

To exploit the forward bias, we can use the regression to find a value for the expected return on a forward position, just like in Equation (7.14). If the expected return is positive (negative), the strategy goes long (short) the foreign currency. While some professional currency managers likely follow such quantitative strategies, deviations from unbiasedness made a much less sophisticated trade popular, namely the **carry trade**.

The idea is simple: Borrow in low-yield currencies such as the yen, and invest in high-yield currencies such as the Australian dollar. The strategy is called “carry” as the carry represents the interest rate differential between the high- and the low-yield currencies. If the exchange rate does not change in value, the investor simply earns the carry. An equivalent strategy is to go long currencies trading at a discount and go short currencies trading at a premium. Again, the naïve idea is that the investor earns the forward discount (the carry) if the future spot rate happens to equal the current spot rate.

Example 7.3 A Carry Trade

Suppose Mrs. Watanabe in Japan faces a spot exchange rate of ¥100/\$ and a 3-month forward rate of ¥99.17/\$. The dollar is trading at an annualized discount in the forward market of

$$4 \times \frac{99.17 - 100}{100} = -3.32\%$$

From covered interest rate parity, we know that this is approximately the interest rate differential between 3-month yen and dollar external currency market investments.

Because the dollar is cheaper in the forward market, Mrs. Watanabe simply buys dollars forward, hoping the spot exchange rate will not change very much. Her eventual return can be decomposed as follows:

$$\text{fmr}(t+1) = s(t+1) - fp(t) = s(t+1) + \frac{3.32\%}{4}$$

The forward discount or carry of $\frac{3.32\%}{4}$ gives her an 83-basis-point cushion. As long as the dollar does not depreciate by more than 83 basis points over the course of the next 3 months, Mrs. Watanabe comes out ahead. Of course, the unbiasedness hypothesis holds that the dollar should be expected to depreciate by exactly 83 basis points!

The carry trade cannot work if the unbiasedness hypothesis holds. Yet, the strategy is different from exploiting the information in the regressions we ran, as it entirely ignores the information in the constant term (see our discussion in the previous subsection). Exploiting the forward bias as implied by regressions makes you primarily invest in currencies where the discount is unusually large relative to historical data, whereas the carry trade simply invests in currencies with high forward discounts (or high interest rates) relative to other currencies.

In Chapter 2, we reported that professional investment firms (such as hedge funds) account for an increasingly larger share of currency market volumes. Over the past decade, a number of hedge funds and other professional investors have started to view investing in currencies as an asset class in its own right. One of the most popular strategies among such investors is the carry trade. Galati et al. (2007) document how carry trade activity increased in the first decade of the 21st century. They also suggest that it may potentially affect currency values by putting upward (downward) pressure on high-yield (funding) currencies and may raise concerns of financial instability, should the carry trade suddenly “unwind,” that is, should the low-interest currencies actually suddenly appreciate.

The carry trade is now viewed as one of the standard currency strategies. For example, in 2006, Deutsche Bank created a carry trade index, easily investable for all types of investors, including retail investors, at a fixed fee. Deutsche Bank’s strategy involves making a diversified investment in equally weighted long or short positions in 10 possible currencies versus the U.S. dollar. The 10 currencies are the euro and the currencies of Australia, Canada, Denmark, Great Britain, Japan, New Zealand, Norway, Sweden, and Switzerland. The strategy involves going long in the three currencies that trade at the steepest forward discounts versus the U.S. dollar (that is, currencies traded in countries where money market yields are higher than those in the United States) and going short in the three currencies that trade at the highest forward premiums versus the U.S. dollar (that is, currencies traded in countries where money market yields are lower than those in the United States). The long or short positions are determined at the beginning of each month and are closed at the end of each month.

Have carry trades been profitable? To judge the profitability of trading strategies, we must introduce some important financial jargon.

Households as Carry Traders?

While you may think that the carry trade is best reserved for professional currency investors, we already pinpointed Mrs. Watanabe in Japan as a retail investor often engaging in carry trades. She is not the only retail investor practicing the carry trade. In Eastern Europe, many households, likely unwittingly, have turned their mortgages into carry trades. Because interest rates in Hungary and Poland were much higher than interest rates on the euro, and especially the Swiss franc, financial institutions started offering mortgages and other loans expressed in foreign currency, mostly Swiss franc. Central bank data reveal that over 50% of mortgages in Hungary are expressed in Swiss francs! The practice is also widespread in Austria, where 13% of households hold Swiss franc-denominated

mortgages, even though the interest differential with the euro is not very large (see Beer et al., 2010). The authors of this study mention that the Austrian households taking out such loans are richer and may be more financially literate than average households. Yet, it is very doubtful that an average household fully understands the risks involved. While they may experience substantial savings on interest costs in the short run, any appreciation of the foreign currency increases the loan amount to be paid off. These risks were painfully realized in Hungary during the first half of 2010, when the forint experienced a 15% depreciation relative to the Swiss franc and, at the same time, Hungarian house prices fell.

Sharpe Ratios and Leverage

To judge the usefulness of a trading strategy, we can compute the economic profits or returns it generates. Because different strategies may have different risks, it is customary to compare the Sharpe ratios of various investment strategies. The Sharpe ratio essentially represents the excess return per unit of volatility. Correcting for volatility is especially important for currency strategies, as they often employ “leverage.” The following analysis reviews the important concepts of leverage and the Sharpe ratio.

The Return on Capital at Risk and Leverage

An investor has a particular amount of capital available to invest, and ultimately we are interested in the return on that capital. However, a forward contract does not necessitate an upfront investment because it is just a bilateral contract with a bank, which means the investor can put more capital at risk than she owns. Because banks want to know that their counterparties can deliver on the contracts, the actual trading strategy typically is to invest the available capital in relatively riskless securities, such as Treasury bills, to absorb potential losses, and then invest possible gains.

If there is exactly \$1 invested in a Treasury bill for every dollar bought or sold in the forward foreign exchange market, the excess return on the trading strategy, that is, the return over and above the return on the Treasury bill, equals the return on “capital at risk.” If forward contracts pertain to more dollars than there are in a

riskless account, the trading strategy uses **leverage**. For example, if for every \$1 in the riskless account, \$2 of forward contracts are made, the leverage ratio is 100%:

$$\begin{aligned}\text{Leverage} &= \frac{\text{Capital at risk} - \text{Capital owned}}{\text{Capital owned}} \\ &= \frac{\$2 - \$1}{\$1} = 100\%\end{aligned}$$

Using leverage in a trading strategy scales up both its returns and its risk. Leverage implies that we should focus on the risk–return trade-off when investigating the profitability of trading strategies. The most popular measure is the **Sharpe ratio**, named after Nobel laureate William F. Sharpe:

$$\text{Sharpe ratio} = \frac{\text{Average excess return}}{\text{Standard deviation of excess return}}$$

Currency Strategies in Practice

The Sharpe ratio in the U.S. stock market is often estimated to be 0.30 to 0.40, meaning that the average annualized excess return is between 5% and 6% and the annualized standard deviation is 15%. Studies find that regression-based foreign exchange strategies produce Sharpe ratios similar and even higher than those available in stock markets, offering a reason for the increase in professional currency managers noted earlier. Bekaert (2011) reports that the assets under management reflected in the Barclay Currency Trader Index (BCTI), an index tracking currency funds, grew from under \$5 billion to over \$25 billion between 2000 and the end of 2007. Pojarliev and Levich (2008) report that the returns and Sharpe ratios on the BCTI initially were quite attractive but have tended to diminish over time, especially over the last few years of the 2000s. However, they identify several currency managers who produced returns with very attractive Sharpe ratios and also outperformed naïve currency strategies, such as the carry trade index.

Although these results are interesting, it is important to realize that past performance need not repeat itself and that currency investing is risky. In particular, in Chapter 3, we indicated that the distribution of currency changes exhibits “fat tails”; that is, extreme outcomes (both positive and negative) are more likely than a normal distribution predicts. If a currency strategy’s return exhibits fat tails, the Sharpe ratio might not adequately reflect the risk–return trade-off.

The global crisis in 2008 proved a wake-up call for the abnormal risks embedded in the carry trade. The Deutsche Bank index performed abysmally, losing more than 20% of its value. This means that a currency fund with a 3-to-1 leverage ratio would have generated a negative return of –80%; in other words, it would have been essentially wiped out. Not surprisingly, many currency funds closed in 2008. Moreover, daily returns on the carry trade index during 2008 were extremely highly correlated with stock returns, suggesting that carry trades do suffer from systematic risk exposure. However, 2008 was not the first time that the carry trade experienced a quick and dramatic unwind. The strategy suffered large losses during the Asian financial crisis of 1997, and again in 1998 when Russia roiled international financial markets by defaulting on its debt in August, the hedge fund Long Term Capital Management collapsed in September, and the yen appreciated very sharply in October. The events in 2008 rekindled interest in two alternative explanations of the forward bias and carry trade returns: risk premiums and peso problems.

Risk Premiums

In the discussion of risk premiums earlier in this chapter, we noted that there are good theoretical reasons that the unbiasedness hypothesis may not hold. Nevertheless, the estimated slope coefficients are quite far from the values implied by the unbiasedness hypothesis. In fact, the regression results imply risk premiums on foreign currency investments must be large and more volatile than expected rates of appreciation, as we show in an advanced section.

Let’s illustrate the ideas with a numerical example. Let the forward discount on the pound relative to the dollar be 2%. However, a bank believes that the pound is expected to appreciate by 3%. What risk premium does the bank expect to earn from investing in pounds? The risk premium is

$$rp(t) = E_t[fmr(t+1)] = E_t[s(t+1) - fp(t)] = 3\% - (-2\%) = 5\%$$

Note that the risk premium is larger than both the expected rate of appreciation and the forward discount. For this forecast to be consistent with a risk explanation, we must believe that the pound is so risky that it not only offers a 2% interest rate premium but also is expected to appreciate by 3%, so that in total, it offers a 5% expected excess return to investors. Is this plausible? We end this section by briefly summarizing the academic debate on whether risk drives the “forward bias.”

The Variability of the Risk Premium¹⁰ (Advanced)

The volatilities of forward premiums on the major currencies are about 3% (on an annualized basis). It turns out that the regression evidence presented in Exhibit 7.5 implies that both the volatilities of expected exchange rate changes and risk premiums are often (much) larger than the volatilities of forward premiums. Let's see why.

The regression states that

$$E_t(s_{t+1}) = a + bfp_t$$

The variance of expected exchange rate changes is therefore

$$\text{VAR}[E_t[s(t+1)]] = \text{VAR}[a + bfp(t)] = b^2\text{VAR}[fp(t)]$$

Hence, if $b^2 > 1$, which is the case for all pairs involving the yen and the \$/£ pair, expected exchange rate changes are more variable than forward premiums. To find the variance of the risk premium, recall that the risk premium is simply the expected forward market return. Therefore,

$$rp(t) = E_t[\text{fmr}(t+1)] = E_t[s(t+1)] - fp(t) = a + (b - 1)fp(t)$$

Hence,

$$\text{VAR}[rp(t)] = (b - 1)^2 \text{VAR}[fp(t)]$$

Consequently, as long as b is negative, which is the case for all currencies, the implied variance of the risk premium is not only larger than the variance of the forward premium, but it is also larger than the implied variance of the expected exchange rate changes.

Is It Risk?

If risk premiums are more variable than expected currency appreciation, a particular movement in the interest rate may more likely be driven by a change in the risk premium than by a change in the expected rate of appreciation of the currency. This is counterintuitive to most economists, who think that most of the forward premium variation reflects expected currency depreciation.

A number of economists (see Frankel and Froot, 1990; and Chinn and Frankel, 2002) have argued that survey data on forecasts of rates of appreciation from market professionals are closely related to forward premiums. The survey data are therefore biased forecasts of rates of appreciation, and the researchers say this indicates that market participants are irrational. There are, however, multiple problems with survey data. Survey participants may not have the proper incentive to tell the truth. In addition, faced with a disparity of forecasts, a statistician must choose something that represents the "market's forecast." Typically, the median forecast is chosen. Ideally, however, we are interested in the marginal investor's expectation. Why is the median of the survey's responses an indication of the opinion of the marginal investor? This calls into question the representativeness of the surveys analyzed in these academic studies.

Nevertheless, basic formal models of risk, such as the CAPM, have a hard time generating risk premiums as variable as implied by the regressions (see, for example, Bekaert, 1996; and Giovannini and Jorion, 1989). The recent global crisis has rekindled interest in the dynamics and economic sources of carry trade returns. The carry trade appears to have attractive long-run returns that trickle in slowly as the "carry" more than compensates for the depreciation of the high-yield currencies. Occasionally, though, a sudden and steep carry trade unwind happens, where the low-yield currencies appreciate sharply, exposing carry traders

¹⁰Fama (1984) was the first to recognize that the estimated slope coefficients in tests of the unbiasedness hypothesis can be interpreted to provide information about the variability of risk premiums and of expected rates of appreciation.

to big losses. Thus, it is said that the carry trade appears to pick up nickels in front of a bulldozer. Statistically, this means the strategy's returns are not normally distributed but exhibit fat tails and negative skewness. Most investors obviously dislike such return properties, and they are not adequately captured by the Sharpe ratio.

Recent academic studies focus on these dynamic properties of carry trade returns to provide new risk-based explanations. Unwinds of the carry trade tend to happen at bad economic times, and it is conceivable that people become dramatically more risk averse when they might lose their job or face large investment losses. Because the returns to carry trades are correlated with such macroeconomic risks, they command a positive risk premium [see Verdelhan (2010) for a recent example of such a model]. Other research focuses on the behavior of traders. Brunnermeier et al. (2009) stress that when a carry trade unwind happens, investment managers face margin calls and may have difficulty funding their levered positions. Their clients may withdraw money as well. These forces cause the managers to unwind their positions, selling the high-yield currencies and buying the low-yield currencies, and in doing so, they exacerbate the losses on the carry trade. If the unwind is bad enough, the investment managers may go out of business. Knowing that this might happen causes an insufficient allocation of risky capital to the carry trade, keeping the returns higher than they should be. This explanation combines the presence of risk premiums with the idea of limits to arbitrage we encountered before.

The new explanations also rely on the fact that there are infrequent disastrous returns to the carry trade. These events by themselves can provide a potential explanation of the forward bias, as we now discuss.

Problems Interpreting the Statistics

Unstable Coefficients in the Unbiasedness Hypothesis Regressions

Exhibit 7.7 presents rolling estimates of the slope coefficients from Equation (7.13) to characterize its dynamics. The first estimate uses the first 5 years of monthly data. The next estimate results from rolling the data forward by 1 month and re-estimating the regression, again with 5 years of data.

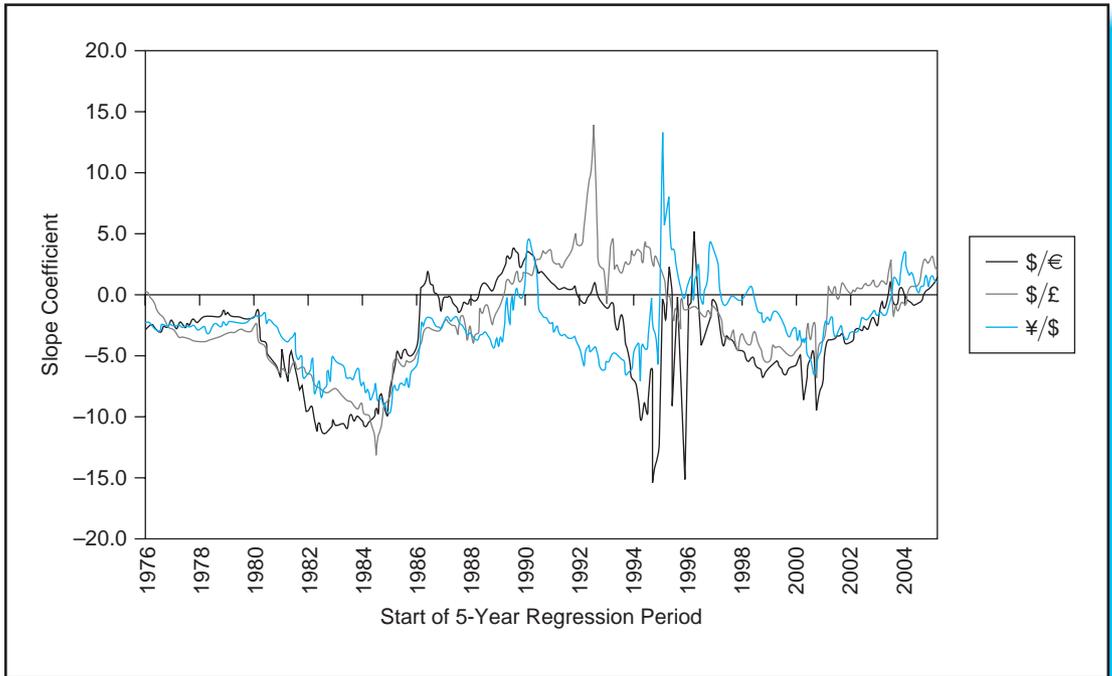
In the regression analysis of the unbiasedness hypothesis, the estimates of the slope coefficient, b , are very far from 1, but Exhibit 7.7 indicates that there is dramatic instability in these coefficients across 5-year intervals. During the major appreciation of the dollar relative to the other major currencies in the early 1980s, the estimated slope coefficient decreased from -5 to -10 . Clearly, this was probably because of the unexpectedly strong appreciation of the dollar and not a response to an increase in the variability of risk premiums. The large carry trade unwinds in the 2007 to 2008 period increased the coefficients towards 1. This evidence indicates a potential problem with the assumption of rational expectations underlying the statistical analysis. We next explain how this might happen.

Peso Problems

A phenomenon called the **peso problem** arises when rational investors anticipate events, typically dramatic, that do not occur during the sample or at least do not occur with the frequency that investors expect. Peso problems invalidate statistical inference conducted under the rational expectations assumption based on data drawn from the period.

The peso problem got its name from considering problems that would have arisen in analyzing Mexico's experience with fixed exchange rates. During 1955 to 1975, the Mexican authorities successfully pegged the peso-dollar exchange rate at MXP12/USD. Suppose we assume that the market sets the forward rate in such a way that it is an unbiased predictor of the future spot rate—that is, we assume that the unbiasedness hypothesis holds. Now, let's see if a statistician would conclude that the forward rate is an unbiased or a biased predictor using the Mexican data.

Exhibit 7.7 Rolling Monthly 5-Year Regression: Monthly Spot Rate Percentage Change Versus Monthly Forward Premium, February 1976–April 2010



Let S_{peg} be the peso–dollar exchange rate at which the Mexican authorities are currently pegging. Let $S_{\text{dev}} > S_{\text{peg}}$ be the rate that the Mexican authorities will choose if they devalue the peso. Suppose that the market knows S_{dev} , and let $\text{prob}(t)$ be the probability that the market assigns to the event that the peso will be devalued during the next month. Then, the 1-month forward rate is an unbiased predictor of the future spot rate when it is the probability-weighted average of the two possible events:

$$F(t) = E_t[S(t+1)] = (1 - \text{prob}(t))S_{\text{peg}} + \text{prob}(t) S_{\text{dev}}$$

The forward rate is the probability of no devaluation multiplied by the current exchange rate plus the probability of a devaluation multiplied by the new exchange rate. As the market’s assessment of the strength of the government’s commitment to the peg changes over time, $\text{prob}(t)$ will change, and so will the forward rate. As long as the devaluation does not materialize, the dollar will trade at a forward premium relative to the peso (in pesos per dollar, $F > S_{\text{peg}}$), and peso money market investments will carry higher interest rates than dollar investments.

Suppose the Mexican authorities successfully peg the peso to the dollar between time T_0 and time T_2 , when they eventually devalue the peso. Suppose also that the market knew during the time period between T_0 and T_2 that the Mexican authorities might devalue the peso at any time. If the statistician takes data from an interval of time during which no devaluation occurs, say, between T_0 and T_1 , where $T_1 < T_2$, and compares forward rates with realized future spot rates, she will conclude that the forward rate is a biased predictor of the future spot rate. During the statistician’s sample, the realized future spot rate is always below the forward rate. Hence, the statistician rejects the null hypothesis that the forward rate is an unbiased predictor of the future spot rate. The statistician has rejected the null hypothesis, but the null hypothesis is true.

How did the statistician go wrong? In other words, what led to the peso problem in this case? When we do statistical analysis on a financial time series using the rational expectations

assumption, we assume that a reasonably long sample of returns is representative of the true distribution of returns that investors thought they faced when they made their investments. For the forward market example, we would assume that the *ex post* spot rates reflect all the possible events that investors thought might happen when they entered into their forward contracts. If there are important events that investors thought might happen but that did not happen, or if relatively rare events happened too frequently, the historical sample means, variances, and correlations in the data may tell us very little about the means, variances, and correlations of returns that investors thought they faced. The historical means, variances, and correlations may also be relatively uninformative about the moments that investors will face in the future. It is in this sense that the past performance of foreign investments may be poor indicators of the returns that investors can expect in the future.

In the case of the Mexican peso, even though the forward rate seemed to be a biased predictor of the future spot rate over 20 years, the devaluation eventually occurred in 1976, thereby validating the prediction embedded in the forward rate.

The Peso Problem and Carry Trades

For the peso problem to explain the evidence regarding carry trade returns and the forward bias we discussed before, the peso events must be anticipated by market participants and, when they occur, they should wipe out the gains accrued before so that excess returns from currency speculation average out to zero. Burnside et al. (2011) claim that even the 2008 disastrous returns do not suffice to make this true. They argue that carry traders can hedge the downside risk using options without sacrificing all their returns, which is inconsistent with a strict interpretation of the peso problem. However, they can explain the carry trade returns if they assume agents become very risk averse when an unwind happens. It appears that time-varying risk premiums remain critical to explain speculative currency returns.

Swedish Interest Rates of 500%

During currency crises, short-term interest rates often become exorbitantly high while long-term interest rates increase only a little, which means there is a large inversion of the term structure of interest rates. This peculiar pattern occurred in Sweden at the height of a currency crisis in Europe in 1992. The Riksbank, Sweden's central bank, raised its marginal lending rate on overnight borrowing to a staggering 500% p.a.—its highest level ever. The marginal lending rate is the rate that applies to the “last resort” financing offered by the Riksbank to Swedish financial institutions when other sources of overnight liquidity have dried up. The marginal lending rate typically provides a ceiling for the overnight market interest rate. Although only a small fraction of the Riksbank's borrowers had to pay the high rate, it still caused the *average* bank borrowing rate to rise to 38%. While interest rates rose on securities of all maturities, the term structure became sharply *inverted*, with 3-month treasury bills yielding 35% and 6-month bills yielding 30%.

Does an interest rate of 500% p.a. make any sense at all? In fact, imposing high interest rates is a tactic that central banks have used successfully since Premier Raymond Poincaré first used it in France in 1924 to prevent speculation against the franc. (This event came to be called “Poincaré's Bear Squeeze.”) With the high borrowing rate, the Swedish government made speculation against the krona prohibitively expensive. It turns out that we can fully understand these interest rate hikes if we use our theory of uncovered interest rate parity and the idea behind the peso problem.

Although the Swedish krona was pegged against the ECU, let us assume for simplicity that it was pegged against the DEM (which had by far the largest weight in the ECU basket). A large fraction of the higher krona interest rates can be accounted for by what is often called a **devaluation premium**—that is, an interest rate that reflects the expected depreciation of a currency. Furthermore, devaluation premiums can also explain the inverted yield curve.

Let's revisit our simple model for exchange rate expectations. For the Swedish krona, there are two possible events:

1. A devaluation with probability of occurrence equal to *prob*
2. No devaluation with probability of occurrence equal to $(1 - \text{prob})$

When the Swedish central bank successfully holds the peg, the exchange rate remains equal to the current spot rate. Let *Z%* denote the magnitude in percentage terms of a devaluation of the krona versus the DEM if the pegged exchange rate does not hold. Then, interest rate differentials tell us something about the probability of devaluation, *prob*, and the percentage magnitude of the devaluation, *Z%*. Consider the expected returns in Swedish krona on two investments for a period of *n* days, with interest rates measured at annual rates and with exchange rates measured in Swedish krona per Deutsche mark as follows:

$$\begin{aligned} \text{Krona investment: } & 1 + i(\text{SKR}) \frac{n}{360} \\ \text{DEM investment: } & \frac{\left[1 + i(\text{DEM}) \frac{n}{360} \right] \times E_t[S(t+n)]}{S(t)} \end{aligned}$$

According to uncovered interest rate parity, these two investments yield the same expected return. Because there are two possible events for the krona—a devaluation or no devaluation—the expected spot rate is simply

$$E_t[S(t+n)] = (1 - \text{prob}) \times S(t) + \text{prob} \times S(t) \times (1 + Z\%)$$

Therefore, by equating the two rates of return, substituting for the expected spot rate, and solving for the intensity of the devaluation (which is the probability of the devaluation multiplied by the size of the devaluation), we find

$$\text{prob} \times Z\% = \frac{1 + i(\text{SEK}) \frac{n}{360}}{1 + i(\text{DEM}) \frac{n}{360}} - 1$$

or by placing the right-hand side over a common denominator, we find

$$\text{prob} \times Z\% = \frac{\left[i(\text{SEK}) \frac{n}{360} \right] - \left[i(\text{DEM}) \frac{n}{360} \right]}{1 + i(\text{DEM}) \frac{n}{360}}$$

Consequently, if krona interest rates are higher than Deutsche mark interest rates, there is a chance of a devaluation of some magnitude. The higher the interest differential, the higher the market assesses the chance and/or the magnitude of a devaluation.

Now, suppose at the height of a currency crisis, *prob* (the likelihood of a devaluation) is very close to 1, say, 0.8. Speculators are quite confident the currency will be devalued, but they are not absolutely sure it will be. Consequently, the interest rate differentials can be used to infer the expected percentage magnitude of the currency devaluation:

	<i>i</i> (SEK)	<i>i</i> (DEM)	<i>prob</i> × <i>Z%</i>	<i>Z%</i> , if <i>prob</i> = 0.8
1 Month	35%	4%	2.57%	3.22%
3 Months	20%	4.5%	3.83%	4.79%

These numbers do not look unreasonable at all.

Why do devaluation expectations of a few percentage points lead to such high interest rates, and why is the effect so much larger for short maturities than for long ones? The inverted yield curve and the large magnitude of the short interest rates are simply a consequence of annualizing interest rates. To make this concrete, suppose that international investors expect a 5% devaluation within a week. Whatever Swedish money market investments they hold, they face an imminent capital loss of 5%. Investors will consequently demand higher interest rates to protect themselves against this possibility. If the interest rate applies to a 1-year maturity, this interest rate increase will be approximately 5%. But when the investment is very short term (such as 1 week), an extra 5% p.a. only means a small increase in the actual return. This won't compensate investors for the capital losses they will suffer as a result of a devaluation. Let the probability of a devaluation be 0.8, and let the DEM interest rate be 3% at the weekly horizon and 5% at the annual horizon. Whatever the investment, $\text{prob} \times Z = 0.8 \times 5\% \equiv 4\%$. According to the formula, we have:

$$\begin{aligned} \text{Devaluation premium} &= \text{1-week investment} = \text{1-year investment} \\ 4\% &= \frac{i(\text{SEK, 1 week})\frac{7}{360} - 3\%\frac{7}{360}}{1 + 3\%\frac{7}{360}} = \frac{i(\text{SEK, 1 year}) - 5\%}{1 + 5\%} \end{aligned}$$

Hence, $i(\text{SEK, 1 week})$ will have to increase by much more than $i(\text{SEK, 1 year})$ to compensate for the expected devaluation of 4%. In particular, we can solve for $i(\text{SEK, 1 week}) = 208.83\%$ p.a., and $i(\text{SEK, 1 year}) = 9.20\%$ p.a. Clearly, the yield curve would be very inverted in this case.

7.7 SUMMARY

This chapter analyzes speculative currency investments. Its main points are as follows:

1. Speculators in currency markets can either borrow currencies they think will weaken while lending currencies they think will strengthen or buy the strengthening currency in the forward market. Speculative currency strategies are only successful when the currency predicted to weaken actually weakens more than the forward rate predicts.
2. Exchange rates are asset prices and are therefore difficult to forecast.
3. The expected return and volatility of a speculative currency investment depend on the mean and the standard deviation, respectively, of the conditional distribution of the future spot exchange rate.
4. Uncovered interest rate parity states that the expected return on an unhedged investment of domestic currency in the foreign money market equals the domestic money market return.
5. The unbiasedness hypothesis states that the forward rate equals the expected future spot rate—that is, what the market expects the spot rate to be on the day your forward contract comes due, $F(t) = E_t[S(t+1)]$. The average forecast error of an unbiased predictor is zero when the average is computed over a large enough sample of forecasts.
6. Both uncovered interest rate parity and the unbiasedness hypothesis are consistent with a narrow view of market efficiency—that is, that there is no expected return to currency speculation. A broader view of market efficiency maintains that the expected profits from a trading strategy should merely compensate the investor for the risk she has taken.
7. The capital asset pricing model (CAPM) provides a theoretical reason why forward rates would be biased predictors of future spot rates and yet the market would still be considered to be efficient. The bias would be attributable to a risk premium, arising from the correlation between forward market returns and the market portfolio return.
8. Whether uncovered interest rate parity and the unbiasedness hypothesis hold has important implications for portfolio management, exchange rate forecasting, and theories of exchange rate determination.

9. If the expected future spot exchange rate and the forward rate differ, hedging transaction exchange risk produces a different revenue or cost than that expected to occur without hedging.
10. If investors have rational expectations, they do not make systematic mistakes when forecasting exchange rates. The actual future rate of appreciation then equals its conditional expectation plus an error term that has a conditional mean of 0; that is, only news makes future exchange rates different from their expected values.
11. The weakest implication of the unbiasedness hypothesis is that the unconditional mean of the forward premium should equal the unconditional mean of the realized rate of appreciation. The data appear consistent with the fact that high-interest-rate or forward-discount currencies tend to depreciate relative to low-interest-rate or forward-premium currencies.
12. Regression tests of the unbiasedness hypothesis indicate that it is strongly inconsistent with the data: Slope coefficients in regressions of the *ex post* rate of appreciation on the forward premium are negative rather than equal to 1. This implies that the forward rate is a biased predictor of the future spot rate.
13. The carry trade goes long in high-yield currencies selling at a forward discount and goes short in low-yield currencies selling at a forward premium.
14. Exploiting the forward bias and carry trades has offered attractive historical returns and Sharpe ratios. These returns may reflect market inefficiency, a risk premium, or a peso problem.
15. A peso problem arises when rational investors anticipate events that do not occur during the sample, or at least not do not occur with the frequency they expect. In such a situation, statistical analysis of returns can be badly biased.
16. In fixed-rate regimes, interest rate differentials provide information about the intensity of a devaluation—that is, the probability of the devaluation multiplied by its magnitude.

QUESTIONS

1. What are two ways to speculate in the currency markets without investing any money up front?
2. What do financial economists mean when they discuss the conditional expectation of the future spot exchange rate?
3. What is the main determinant of the variability of forward market returns?
4. Describe how you construct the uncertain yen-denominated return from investing 1 yen in the Swiss franc money market.
5. What is a hedged foreign currency investment? What happens if you hedge your return in Question 4?
6. What does it mean for the 90-day forward exchange rate to be an unbiased predictor of the future spot exchange rate?
7. Why is it true that the hypothesis that the forward exchange rate is an unbiased predictor of the future spot exchange rate is equivalent to the hypothesis that the forward premium (or discount) on a foreign currency is an unbiased predictor of the rate of its appreciation (or depreciation)?
8. It is often claimed that the forward exchange rate is set by arbitrage to satisfy (covered) interest rate parity. Explain how interest rate parity can be satisfied and how the forward exchange rate can be set by speculators in reference to the expected future spot exchange rate.
9. It is sometimes asserted that investors who hedge their foreign currency bond or stock returns remove the foreign exchange risk associated with the investment, reduce the volatility of their domestic currency returns, and thus get a “free lunch” because the mean return in domestic currency remains the same as the mean return in the foreign currency. Is this true or false? Why?
10. It is often argued that forward exchange rates should be unbiased predictors of future spot exchange rates if the foreign exchange market is efficient. Is this true or false? Why?
11. What is the prediction of the CAPM for the relationship between the forward exchange rate and the expected future spot exchange rate?
12. If the CAPM explains deviations of the forward exchange rate from the expected future spot exchange rate, explain why one party involved in a forward contract would be willing to enter into a contract with an expected loss.
13. Why is it only the covariance of an asset’s return with the return on the world market portfolio that determines whether there is a risk premium associated with the asset’s expected return?
14. What is the rational expectations hypothesis, and how is it applied to tests of hypotheses about expected returns in financial markets?

15. Suppose that the forward premium equals the conditional expectation of the future rate of appreciation of the foreign currency relative to the domestic currency. If we form the average realized rate of appreciation from a large sample of data and compare it to the average forward premium, what should be true?
16. Explain how you would use a regression to test the unbiasedness hypothesis.
17. Suppose you regress the realized rate of appreciation of a foreign currency on a constant and the forward premium on the foreign currency. What interpretation can you give to the estimated slope coefficient? If the slope coefficient is negative, is it true that the forward premium is predicting the wrong sign for the rate of appreciation?
18. What does a negative slope coefficient in an unbiasedness regression imply about the variability of risk premiums relative to variability of expected rates of appreciation?
19. What is a carry trade?
20. What is a Sharpe ratio?
21. Do carry trades contain risks that may not be reflected in their Sharpe ratios?
22. What is a peso problem? Explain the term within the context of its original derivation. Now, explain how peso problems can generally plague the study of financial market returns.
23. How can you use interest rate differentials to understand the probability of a devaluation and the potential magnitude of the devaluation?

PROBLEMS

1. Over the next 30 days, economists forecast that the pound may weaken relative to the dollar by as much as 7%, or it may strengthen by as much as 6%. The possible rates of change are $-7%$, $-5%$, $-3%$, $-1%$, $0%$, $2%$, $4%$, and $6%$. If these values are equally likely, what are the mean and standard deviation of the future spot exchange rate if the current rate is $\$1.5845/\pounds$?
2. Consider the following hypothetical facts about Mexico: The peso recently lost over 40% of its value relative to the dollar. Over the course of the next 90 days, there is a 35% chance that the Mexican government will lose control of the economy. If it does, the peso will lose 33% of its value relative to the dollar, and the Mexican stock market will fall by 39%. Alternatively, the U.S. Congress may vote to help Mexico by offering collateral for Mexican government loans. In that case, the peso will appreciate 27% relative to the dollar, and the Mexican stock market will rise by 29%. As a U.S. investor with no current assets or liabilities in Mexico, you have decided to speculate. Calculate your expected dollar return from investing dollars in the Mexican stock market for the next 90 days.
3. Suppose that the 90-day forward rate is $\$1.19/\pounds$, the current spot rate is $\$1.20/\pounds$, and you expect the future spot rate in 90 days to be $\$1.21/\pounds$. What contract would you make to speculate in the forward market by either buying or selling $\pounds 10,000,000$? What is your expected profit? If the standard deviation of the 90-day rate of appreciation of the euro relative to the dollar is 3%, what range covers 95% of your possible profits and losses?
4. Suppose the rate of appreciation of the dollar relative to the yen over the next 90 days is normally distributed with a mean of $-1%$ and a standard deviation of 3%. Use a spreadsheet program to graph the distribution of the future yen-dollar exchange rate. If the current spot exchange rate is $\pounds 99/\pounds$, and the 90-day forward rate is $\pounds 98.30/\pounds$, describe the distribution of yen profits or losses from selling $\$5,000,000$ forward?
5. Suppose that the spot exchange rate is $\$1.55/\pounds$, that the beta on a forward contract to buy pounds with dollars is 1.5, and that the expected excess dollar rate of return on the market portfolio is 7%. What is the expected profit or loss on a forward purchase of $\pounds 1,000,000$? Explain how this can be an equilibrium.
6. Suppose the estimated slope coefficient in a regression of the rate of depreciation of the dollar relative to the yen on a constant and the forward discount on the dollar is -2 , and the standard deviation of the forward discount, measured on an annualized basis, is 2.5%. What is a lower bound for the variability of the risk premium in the yen-dollar forward market?
7. Suppose the British pound (GBP) is pegged to the euro (EUR). You think there is a 5% probability that the GBP will be devalued by 10% over the course of the next month. What interest differential would prevent you from speculating by borrowing GBP and lending EUR?

8. Argentina's monetary stabilization plan in 1991 included introducing a currency board that tied the Argentine peso (ARS) to the U.S. dollar at an exchange rate of ARS1/USD1. On June 21, 2000, the 3-month interest rates quoted by Argentine banks were 6.71% in USD and 7.33% in ARS. Suppose the difference reflected some probability that the currency board would be abandoned and the peso devalued, and investors think a 10% devaluation to ARS1.10/USD is possible. What is the probability of this happening if uncovered interest rate parity holds? In early 2001, confidence in the currency board eroded and interest rates soared to well over 10%. What is the possibility of a 10% devaluation if the 3-month interest rates are 20% in ARS and 6.0% in USD?
9. The British bank Barclays has developed an exchange-traded note that pays off the Barclays Capital Intelligent Carry Index™. Look up information on this index on the Web. Explain why you like or dislike Barclays's strategy.

BIBLIOGRAPHY

- Beer, Christian, Steven Ongena, and Marcel Peter, 2010, "Borrowing in Foreign Currency: Austrian Households as Carry Traders," *Journal of Banking and Finance* 34, pp. 2198–2211.
- Bekaert, Geert, 1996, "The Time Variation of Risk and Return in Foreign Exchange Markets: A General Equilibrium," *Review of Financial Studies* 9, pp. 427–470.
- _____, 2011, "Valuing Currency Management: TOM vs. U.S. Commerce Bank," Columbia CaseWorks No. 100310.
- Brunnermeier, Markus, Stefan Nagel, and Lasse Pedersen, 2009, "Carry Trades and Currency Crashes," Chapter 5 in *NBER Macroeconomics Annual 2008*, Cambridge, MA: National Bureau of Economic Research.
- Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo, 2011, "Do Peso Problems Explain the Returns to the Carry Trade?" *Review of Financial Studies*, forthcoming.
- Chinn, Menzie, and Jeffrey A. Frankel, 2002, "Survey Data on Exchange Rate Expectations: More Currencies, More Horizons, More Tests," in Bill Allen and David Dickinson, eds., *Monetary Policy, Capital Flows and Financial Market Developments in the Era of Financial Globalisation: Essays in Honour of Max Fry*, London: Routledge.
- Engel, Charles, 1996, "The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence," *Journal of Empirical Finance* 3, pp. 123–191.
- Fama, Eugene, 1984, "Forward and Spot Exchange Rates," *Journal of Monetary Economics* 14, pp. 319–338.
- Frankel, Jeffrey, and Kenneth Froot, 1990, "Exchange Rate Forecasting Techniques, Survey Data, and Implications for the Foreign Exchange Market," in D. Das, ed., *Current Issues in International Trade and International Finance*, Oxford, U.K.: Oxford University Press.
- Froot, Kenneth A., and Richard H. Thaler, 1990, "Foreign Exchange," *Journal of Economic Perspectives* 4, pp. 179–192.
- Galati, Gabriele, Alexandra Heath, and Patrick McGuire, 2007, "Evidence of Carry Trade Activity," *BIS Quarterly Review* 3, pp. 27–41.
- Giovannini, Alberto, and Philippe Jorion, 1989, "The Time-Variation of Risk and Return in the Foreign-Exchange Stock Markets," *Journal of Finance* 44, pp. 307–325.
- Hansen, Lars Peter, 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica* 50, pp. 1029–1054.
- Jylhä, Petri, Jussi-Pekka Lyytinen, and Matti Suominen, 2010, "Arbitrage Capital and Currency Carry Trade Returns," Helsinki School of Economics Working Paper W-459.
- King, Michael, and Dagfinn Rime, 2010, "The \$4 Trillion Question: What Explains FX Growth Since the 2007 Survey?" *BIS Quarterly Review*, December, pp. 27–42.
- Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2009, "Common Risk Factors in Currency Markets," National Bureau of Economic Research Working Paper W14082.
- Pojarliev, Momtchil, and Richard M. Levich, 2008, "Do Professional Currency Managers Beat the Benchmark?" *Financial Analysts Journal* 64, pp. 18–32.
- Sharpe, William F., 1964, "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk," *Journal of Finance* 19, pp. 425–442.
- Siegel, Jeremy J., 1972, "Risk, Information, and Forward Exchange," *Quarterly Journal of Economics* 86, pp. 303–309.
- Verdelhan, Adrien, 2010, "A Habit-Based Explanation of the Exchange Rate Risk Premium," *Journal of Finance* 65, pp. 123–145.
- Villanueva, O. Miguel, 2007, "Forecasting Currency Excess Returns: Can the Forward Bias Be Exploited?" *Journal of Financial and Quantitative Analysis* 42, pp. 963–990.

The Siegel Paradox

Suppose we consider Blake Bevins, Kevin Anthony's British counterpart, who is investing in the dollar money market. Let $\bar{S}(\text{£}/\$)$ and $\bar{F}(\text{£}/\$)$ denote the pound/dollar spot and forward exchange rates, so at each point of time $\bar{S}(\text{£}/\$) = \frac{1}{S(\$/\text{£})}$. Now, apply Equation (7.5) from the British perspective, $\bar{F}(t) = E_t[\bar{S}(t+1)]$. But, of course, $\bar{F}(\text{£}/\$) = \frac{1}{F(\$/\text{£})}$, so that

$$\begin{aligned} E_t[\bar{S}(t+1, \text{£}/\$)] &= \frac{1}{F(t, \$/\text{£})} \\ &= \frac{1}{E_t[S(t+1, \$/\text{£})]} \end{aligned}$$

So, for the unbiasedness hypothesis to hold from both the British and American perspectives, it must be the case that

$$\begin{aligned} E_t[\bar{S}(t+1, \text{£}/\$)] &= E_t\left[\frac{1}{S(t+1, \$/\text{£})}\right] \\ &= \frac{1}{E_t[S(t+1, \$/\text{£})]} \end{aligned}$$

However, we know the latter equality is false because of a statistical property known as Jensen's inequality.¹¹

Rather than get mired in statistical jargon, let's work out a simple numeric example. Suppose Kevin and Blake agree on the following possible scenarios for the future exchange rate:

	S(t+1, \$/£)	S̄(t+1, £/\$) = $\frac{1}{S(t+1, \$/\text{£})}$	Probability
Scenario 1	1.50	0.6667	0.714
Scenario 2	1.65	0.6061	0.286

From Kevin's perspective, the expected future \$/£ exchange rate is

$$\begin{aligned} E_t[S(t+1, \$/\text{£})] &= 0.714 \times \$1.50/\text{£} \\ &\quad + 0.286 \times \$1.65/\text{£} \\ &= \$1.5429/\text{£} \end{aligned}$$

This is the forward rate derived earlier. According to Blake, the expected £/\$ rate is

$$\begin{aligned} E_t[\bar{S}(t+1, \text{£}/\$)] &= 0.714 \times \text{£}0.6667/\$ \\ &\quad + 0.286 \times \text{£}0.6061/\$ \\ &= \text{£}0.6493/\$ \end{aligned}$$

Is this consistent with the \$1.5429/£ rate? The answer is no because

$$0.6493 \neq \frac{1}{1.5429} = 0.6481$$

We see that when the unbiasedness hypothesis is considered from the two different currency perspectives, it leads to an inconsistency. We cannot have two different forward rates in the market! This little conundrum is known as the Siegel paradox because Jeremy Siegel (1972) was the first to point out this inconsistency.

Whereas some have argued that the Siegel paradox invalidates the unbiasedness hypothesis as a reasonable theory, note that the difference between 0.6481 and 0.6493 is small: In percentage terms, it represents less than a 0.2% difference. Hence, we will ignore the Siegel paradox for the remainder of this book. Moreover, it is possible to formulate versions of the unbiasedness hypothesis either using logarithmic exchange rates or using real values that resolve the Siegel paradox (see Engel, 1996).

¹¹In fact, because $f(x) = \frac{1}{x}$ is a convex function, Jensen's inequality implies $E_t\left[\frac{1}{S(t+1)}\right] > \frac{1}{E_t[S(t+1)]}$.

The Portfolio Diversification Argument and the CAPM

If an investor places all her wealth in only one asset, the asset's expected return and variance are the mean and variance of the investor's portfolio. The purpose of this appendix is to review how the mean and variance of a portfolio are determined when there is more than one asset in the portfolio. To do this easily, we must develop some notation. Let R_i be the return on asset i and denote the expected value or mean return on asset i as $E(R_i)$. Let σ_{ij} denote the covariance between the returns on asset i and asset j . Covariance is a measure of the degree to which two returns move together, and it is found by taking the expectation of the product of the deviations of the returns from their respective means:

$$\sigma_{ij} = E\left[[R_i - E(R_i)][R_j - E(R_j)]\right]$$

Because the covariance involves the product of two random variables and the order of multiplication is unimportant, $\sigma_{ij} = \sigma_{ji}$. Also, from the definition of variance, which is the expected value of the squared deviation around the mean, we have

$$\sigma_{ii} = E\left[[R_i - E(R_i)]^2\right]$$

The square root of the variance is the standard deviation. Often, people find it more intuitive to think in terms of correlations between returns on assets rather than covariances because the correlation is a number between -1 and 1 . The correlation coefficient, ρ_{ij} , is defined to be the covariance divided by the product of the standard deviations of the two assets:

$$\rho_{ij} = \frac{\sigma_{ij}}{\sqrt{\sigma_{ii}}\sqrt{\sigma_{jj}}} \quad (7A.1)$$

Now, we can examine the mean and variance of the return on a portfolio of several assets. Let w_i denote the share of the investor's wealth that is invested in asset i . Let's also begin with just two assets in the portfolio. Suppose the investor puts a share of her wealth equal to w_1 in asset 1 and the remainder of her wealth in asset 2, such that $w_2 = 1 - w_1$.

The actual return on the portfolio, R_p , will be the weighted average of the returns on the two assets, where the weights are the shares of invested wealth:

$$R_p = w_1R_1 + w_2R_2 \quad (7A.2)$$

Hence, to find the mean return on the portfolio, we take the expectation of the realized return in Equation (7A.2), and we find

$$E(R_p) = w_1E(R_1) + w_2E(R_2)$$

Just as the actual return is a weighted average of the actual individual returns, the expected return on the portfolio is the same weighted average of the expected returns on the assets.

The variance of the return on the portfolio $V(R_p)$ is the expectation of the squared deviation of the return from its mean, as in the following:

$$V(R_p) = E\left[\left[(w_1R_1 + w_2R_2) - (w_1E(R_1) + w_2E(R_2))\right]^2\right] \quad (7A.3)$$

By multiplying out and rearranging the terms in Equation (7A.3), we find that

$$\begin{aligned} V(R_p) &= w_1^2E\left[[R_1 - E(R_1)]^2\right] \\ &\quad + w_2^2E\left[[R_2 - E(R_2)]^2\right] \\ &\quad + 2w_1w_2E\left[[R_1 - E(R_1)][R_2 - E(R_2)]\right] \\ V(R_p) &= w_1^2\sigma_{11} + w_2^2\sigma_{22} + 2w_1w_2\sigma_{12} \end{aligned}$$

Let's do a calculation with some real numbers to see how the mean and variance of a portfolio are related to the means and variances of the individual assets. Suppose that the expected return on asset 1 is 9%, and its standard deviation is 22%, whereas the expected return on asset 2 is 11%, and its standard deviation is 24%. Suppose also that the correlation between the returns on the two assets is 0.4, and from Equation (7A.1), we find that the covariance between the two returns is $\sigma_{12} = (0.4)(0.22)(0.24) = 0.02112$.

Now, we can calculate the mean and variance of any portfolio composed of assets 1 and 2. Suppose we put 35% of our wealth in asset 1 and 65% in asset 2. The mean return on our portfolio is then

$$E(R_p) = (0.35)(0.09) + (0.65)(0.11) = 0.1030$$

and the variance of the return on our portfolio is

$$V(R_p) = (0.35)^2(0.22)^2 + (0.65)^2(0.24)^2 + 2(0.35)(0.65)(0.02112) = 0.039875$$

The standard deviation of our portfolio is therefore $\sqrt{0.039875} = 0.1997$ or 19.97%.

The ratio of the mean to the standard deviation of an asset or a portfolio is a measure of the trade-off an investor faces between return and risk. For asset 1, the ratio of mean to the standard deviation is $09\%/22\% = 0.41$, and for asset 2, it is $11\%/24\% = 0.46$. For the portfolio, the ratio of the mean to the standard deviation is $10.30\%/19.97\% = 0.52$. By diversifying across the two assets, we have improved our risk–return trade-off. Also, note that the standard deviation of the portfolio is lower than the standard deviation of either asset. Diversification makes some risk disappear.

Because there are many more than two assets in the world, we next want to examine what happens if we put a small amount of our wealth in each of N assets. To further simplify the analysis, let's put an equal share, $w_i = (1/N)$, in the N different assets. The portfolio's mean return is just the weighted sum of the expected returns on the N assets, as in Equation (7A.2):

$$E(R_p) = \sum_{i=1}^N w_i E(R_i) = \sum_{i=1}^N \frac{E(R_i)}{N}$$

Consequently, the portfolio's mean return is the average of the mean returns on the N assets.

The variance of the return on an N -asset portfolio is as follows:

$$V(R_p) = E \left[\sum_{i=1}^N w_i [R_i - E(R_i)] \sum_{j=1}^N w_j [R_j - E(R_j)] \right] \quad (7A.4)$$

If you multiply out the terms involving the summations on the right-hand side of Equation (7A.4), you will find that you must take the sum of the expectations of N^2 terms. There will be N variances that arise from the multiplication of the return on an asset with itself, and there will be $N(N-1)$ other terms involving covariance terms. So, there will be $N(N-1)/2$ distinct covariance terms because $\sigma_{ij} = \sigma_{ji}$. In Equation (7A.4), the weights are multiplied by each other, but because the weights on the equal-weighted portfolio are the same, each of the N^2 terms in Equation (7A.4) is multiplied by $1/N^2$. Therefore,

$$V(R_p) = \frac{1}{N^2} \sum_{i=1}^N \sigma_{ii} + \frac{2}{N^2} \sum_{i=1}^{N-1} \sum_{j=i+1}^N \sigma_{ij} \quad (7A.5)$$

The double summation term, $\sum_{i=1}^{N-1} \sum_{j=i+1}^N \sigma_{ij}$, is multiplied by 2 because the summation involves only the distinct $N(N-1)/2$ covariances. Let's define the average variance as

$$\Lambda_i = \frac{1}{N} \sum_{i=1}^N \sigma_{ii}$$

and the average covariance as

$$\Lambda_{ij} = \frac{1}{N(N-1)/2} \sum_{i=1}^{N-1} \sum_{j=i+1}^N \sigma_{ij}$$

Equation (7A.5) implies that the portfolio variance can be written as

$$V(R_p) = \frac{1}{N} \Lambda_i + \left(1 - \frac{1}{N}\right) \Lambda_{ij} \quad (7A.6)$$

Notice that as N gets large in Equation (7A.6), the importance of the average variance goes to zero. Thus, as N gets large, the variance of the return on a highly diversified portfolio is driven to be equal to the average covariance of the assets in the portfolio. If asset returns were uncorrelated, the average covariance would be zero, and a highly diversified portfolio would produce an essentially riskless return, even though each of the individual asset returns was itself quite variable. Notice also that assets with negative covariances are very important because they reduce the average covariance of the portfolio.

From Equation (7A.6), it is clear that the individual variance of an asset will not affect the overall variance of the portfolio, and the individual variance consequently should not affect the expected return that a risk-averse investor demands to hold that particular asset. This intuition leads directly to the CAPM as a relationship describing how expected returns are determined. Essentially, the CAPM builds on the intuition that an investor will add an asset to his portfolio until he cannot further improve the risk–return trade-off of the portfolio. We elaborate on this intuition in Chapter 13.

Although the large portfolio in our analysis was arbitrary, the fundamental insight of the CAPM was that with a few additional assumptions, it would have to be the case that, in equilibrium (that is, when all investors are happily holding the existing assets in the marketplace at their current prices, without feeling the need to trade them), the well-diversified portfolio that every investor would hold would be the market portfolio. All investors would hold some fraction of their wealth in the market portfolio, with more risk-averse investors holding smaller fractions and more risk-tolerant investors holding larger fractions and possibly borrowing to invest in the market.

A Regression Refresher

In Section 7.5, we tested the unbiasedness hypothesis with a linear regression model:

$$y(t+1) = a + bx(t) + \varepsilon(t+1)$$

where the dependent (or explained) variable $y(t+1)$, which was the rate of appreciation, $s(t+1)$, is regressed on an independent (or explanatory) variable, $x(t)$, which was the forward premium, $fp(t)$. The regression describes how variation in $y(t+1)$ can be explained linearly by variation in $x(t)$. We want to find values of the parameters, a and b , that make $a + bx(t)$ as close to $y(t+1)$ as possible. The fit is unlikely to be perfect, so there will be an error (or disturbance) term, as indicated by $\varepsilon(t+1)$.

Econometricians have developed several methods to find “estimates,” or values, for the parameters, a and b , given data on $y(t+1)$ and $x(t)$. For any given sample of data, these estimates are just numbers and are typically represented by \hat{a} and \hat{b} . With such estimates, we can compute the actual errors, called *residuals*, that the model makes in predicting $y(t+1)$:

$$\hat{\varepsilon}(t+1) = y(t+1) - \hat{a} - \hat{b}x(t)$$

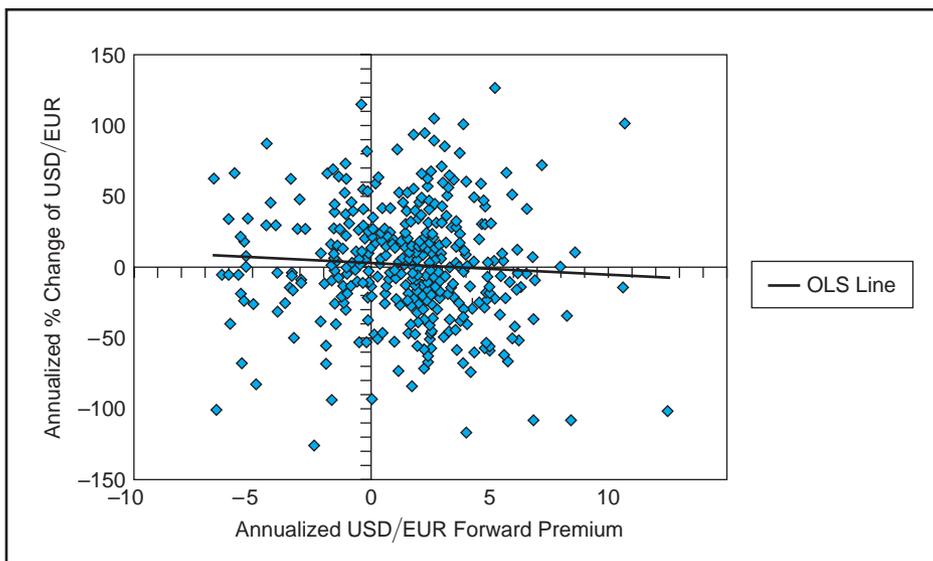
The formula by which the data are transformed into an actual estimate is called an **estimator**, and the most

popular estimator for the linear regression model is the **OLS estimator**. OLS stands for *ordinary least squares* because the estimator minimizes the sum of the squared residuals. That is, the estimates of a and b are such that the sum of the squared residuals, $\sum_{t=1}^T \hat{\varepsilon}(t+1)^2$, is as low as possible, and we are assuming that we have $T+1$ total observations, of which only T will be used in the regression.

To illustrate this concretely, let’s go back to the actual monthly data on dollar/euro exchange rates and forward premiums used for Exhibit 7.5, which were between February 1976, and April 2010. The monthly exchange rate changes represent our $y(t+1)$ observations; the forward premiums represent our $x(t)$ observations. We have to be careful with the timing to match up, say, the April 2001 exchange rate change with the forward premium for the end of March 2001.

Exhibit 7A.1 presents a scatter plot of the data, with the exchange rate changes on the vertical axis and the forward premiums on the horizontal axis. The OLS regression line through this scatter plot minimizes the sum of the squared deviations between the actual data and the regression line. The corresponding fitted values that lie on the regression line are also on the graph.

Exhibit 7A.1 Regression Residuals with Fitted Values



Concretely, the OLS estimator resulting from this procedure for the slope of the line is

$$\hat{b} = \frac{\frac{1}{T} \sum_{t=1}^T [y(t+1) - \bar{y}][x(t) - \bar{x}]}{\frac{1}{T} \sum_{t=1}^T [x(t) - \bar{x}]^2}$$

where $\bar{y} = (1/T) \sum_{t=1}^T y(t+1)$ and $\bar{x} = (1/T) \sum_{t=1}^T x(t)$ are the sample means, and $\hat{a} = \bar{y} - \hat{b}\bar{x}$ is the constant. Note that the numerator of \hat{b} represents an estimate of the covariance between $y(t+1)$ and $x(t)$, whereas the denominator represents an estimate of the variance of $x(t)$. Hence, the slope coefficient b is the covariance of the dependent variable and the independent variable divided by the variance of the independent variable:

$$b = \frac{\text{cov}[y(t+1), x(t)]}{\text{var}[x(t)]}$$

When we carry out the actual regression with the data given in Exhibit 7A.1, we find:

$$\begin{array}{cc} \hat{a} = 3.26 & \hat{b} = -0.84 \\ (2.31) & (0.81) \\ [0.84] & [0.98] \\ R^2 = 0.004\% \end{array}$$

Note that we annualized the constant \hat{a} by multiplying by 12.

An OLS regression also yields a standard error for the estimates, which gives an idea of how confident we

are in the estimates. We report standard errors in parentheses below the parameter estimates as shown in the previous equation; that is, the standard error of \hat{a} is 2.31, for example. Even if $y(t+1)$ and $x(t)$ are totally independent, they may appear to be related just by chance. Use of the standard error together with the coefficient estimate allows computation of a confidence level for b to be different from a particular value. For example, the unbiasedness hypothesis in the context of the regression model represents the null hypothesis $\hat{b} = 1$. We would like to know whether \hat{b} is close to or far away from 1 in a statistical sense.

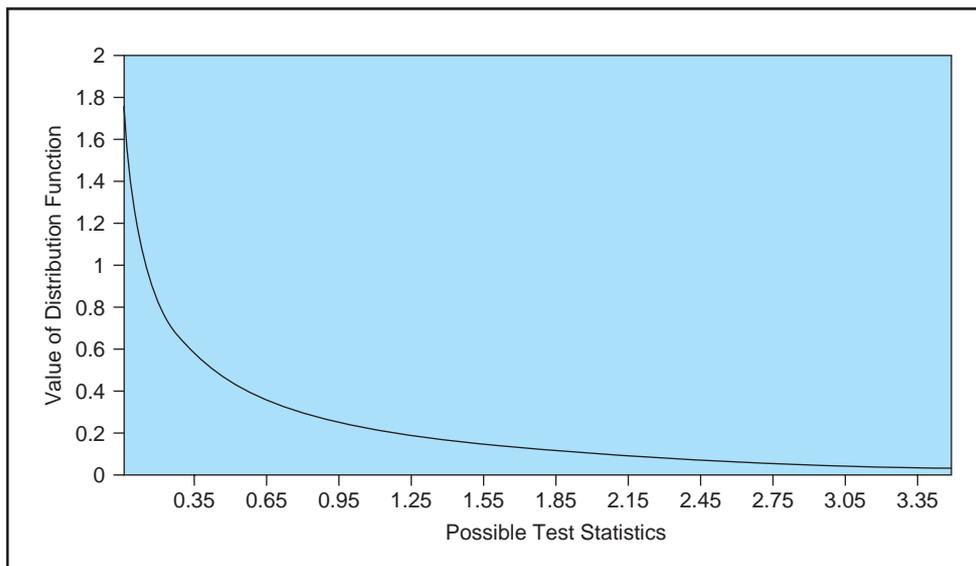
If we want to test whether b is 1, we compute the square of $\hat{b} - 1$ divide by the standard error of \hat{b} . Let us introduce the test statistic z :

$$z = \left[\frac{\hat{b} - 1}{\text{se}(\hat{b})} \right]^2$$

If \hat{b} is truly close to 1, the value of z should be small, and if the true value of b is not equal to 1, the z statistic should be large. However, the true b may be far from 1, but the estimate may be very noisy—that is, the standard error may be big. In this case, our test statistic z will be small as well. In our sample regression, the standard error for \hat{b} is 0.81; hence, $z = 5.1602$. Standard errors are inversely related to the size of the sample, and our sample here is quite long, so that z is relatively large. But at what value of z do we reject the null hypothesis?

If the sample is large, econometricians have actually figured out that the statistic z should follow a particular

Exhibit 7A.2 Chi-Square Distribution



statistical distribution if the null hypothesis is correct. In our case, this distribution is a chi-square distribution with one degree of freedom. Exhibit 7A.2 graphs a $\chi^2(1)$ distribution. Even if the null hypothesis is true, sometimes, by chance, large values of z might occur, but they are not very likely. The higher z is, the less likely it is that z comes from a $\chi^2(1)$ distribution. In fact, only 5% of the observations of $\chi^2(1)$ distribution should be above 3.841. Hence, if our test statistic yields a value higher than 3.841, we are more than 95% confident that the null hypothesis is rejected because there is more than a 95% chance that a $\chi^2(1)$ variable is lower than the z statistic.

Statisticians often focus on “5% level” tests. The value 3.841 is called the critical value of the $\chi^2(1)$ distribution for a 5% test, and when z exceeds the critical value, we say that the null hypothesis is rejected at the 5% level. In the chapter itself, we primarily focus on these confidence levels. In this example, the confidence level is 0.98. We report these confidence levels in square

brackets above. Consequently, we quite confidently reject the hypothesis.

The null hypothesis does not necessarily have to be about just one coefficient. We can also test multiple restrictions together (for instance, $\hat{a} = 0$ and $\hat{b} = 1$), and the resulting statistic will follow a chi-square distribution with degrees of freedom equal to the number of restrictions tested.

Finally, the regression output typically also provides the R^2 statistic. This statistic measures how much of the variation of the dependent variable is explained by the regression model. Concretely, it is computed as the variance of $\hat{a} + \hat{b}x(t)$ divided by the variance of $y(t+1)$. The R^2 is very low in our example because the regression is predictive: We use a variable at time t to predict changes in an asset price at time $t+1$. Most of the variation in the exchange rate will be driven by news that is by definition unpredictable. In Exhibit 7A.1, the poor R^2 is obvious as the data points are often quite far away from the regression line.