CHAPTER 5

Risk Management Strategies: Futures

This chapter builds on Chapters 3 and 4. In Chapter 3, we learned the mechanics of expected return and risk. In Chapter 4, we learned about the no-arbitrage price relation that links the price of a forward/futures to the price of its underlying asset. This chapter explores the role of forward/futures contracts in managing expected return and risk. In moving forward through the chapter, we will use only the term "futures" rather than "forward and futures" for expositional convenience. The risk management techniques apply to both contracts equally well. The decision to use "futures" rather than "forwards" is based on the fact that historical futures data are more broadly available for estimation purposes. Since the chapter deals with expected return and risk, the most natural place to begin is with a demonstration of how futures fit within the capital asset pricing model (CAPM). We then focus on using futures contracts to manage different types of risks. We begin with price risk and show how an airline can hedge the cost of jet fuel. Next we focus on revenue risk and show how a farmer can hedge the sales proceeds of his crop in an environment with both price and quantity risks. For other corporate risk managers, gross margin (i.e., uncertain revenue less uncertain costs) risk is often the primary risk management focus. Oil refiners, for example, are concerned about the difference between the revenue they realize through the sale of heating oil and unleaded gasoline and the cost of the crude oil they must acquire to produce these products. For fund managers, more than one risk factor may be affecting portfolio value. Someone managing a junk bond portfolio, for example, faces both interest rate and stock market risk exposures. We show how to incorporate multiple risk factors in setting the optimal hedge. The chapter concludes with a brief summary.

EXPECTED RETURN AND RISK

Like other risky financial instruments, futures contracts have expected returns and risks that can be modeled within the CAPM. The key to understanding exactly how lies in the relation between the rate of price change of a futures and the rate of return of its underlying asset.

To begin, we recall the net cost of carry relation,

$$F_t = S_t e^{(r-i)(T-t)} \tag{5.1}$$

where the subscript t has been added to denote a particular point in time prior to the contract's expiration. Taking the natural logarithm of both sides of (5.1) provides

$$\ln F_t = (r-i)(T-t) + \ln S_t \tag{5.2}$$

Now consider the transformed net cost of carry relation (5.2) an instant earlier in time at $t + 1$, that is,

$$\ln F_{t-1} = (r-i)(T-t+1) + \ln S_{t-1} \tag{5.3}$$

Subtracting (5.3) from (5.2), we find that the continuous rate price change of the futures is

$$R_F \equiv RA_F \equiv \ln(F_t/F_{t-1}) = -(r-i) + \ln(S_t/S_{t-1}) \tag{5.4}$$

In equation (5.4), R_F denotes the rate of return on a futures contract and RA_F is its rate of price appreciation. We make this distinction to underscore the fact that the only income arising from holding a futures contract is price change.[1] The rate of return from investment in the underlying asset, R_S, on the other hand, is the sum of two components—the continuous rate of price appreciation $RA_S \equiv \ln(S_t/S_{t-1})$ and the income rate i. The relation between the random returns of the futures and its underlying asset is therefore

$$\tilde{R}_F = \tilde{R}_S - r \tag{5.5}$$

where tildes have been added to distinguish between what is uncertain (i.e., the returns on the futures and its underlying asset) from what is certain (i.e., the risk-free rate of interest).

Expected Return-Risk Relation

With the return relation (5.5) in hand, the role of futures contracts within the CAPM is easily uncovered. To do so, first note the expected return on a futures contract equals the expected return on the underlying asset less the risk-free rate of interest, that is,

[1] The distinction also serves to combat the criticism that, since the futures involves no net investment, the rate of return on a futures is undefined.

$$E_F = E_S - r \tag{5.6}$$

Next, note that total risk (as measured by return variance or its square root, standard deviation) and market risk (as measured by beta) of the futures contract equal the total risk and the market risk of the underlying asset. The variance of futures return equals the variance of the asset return,[2]

$$\text{Var}(\tilde{R}_F) = \text{Var}(\tilde{R}_S - r) = \text{Var}(\tilde{R}_S) \tag{5.7}$$

and the beta of the futures contract equals the beta of the underlying asset,

$$\beta_F \equiv \frac{\text{Cov}(\tilde{R}_F, \tilde{R}_M)}{\text{Var}(\tilde{R}_M)} = \frac{\text{Cov}(\tilde{R}_S - r, \tilde{R}_M)}{\text{Var}(\tilde{R}_M)} = \frac{\text{Cov}(\tilde{R}_S, \tilde{R}_M)}{\text{Var}(\tilde{R}_M)} \equiv \beta_S \tag{5.8}$$

Hence, while the risks of the futures contract are the same as those of the underlying asset, the expected return of the futures is below the expected return of the underlying asset by an amount equal to the risk-free rate of interest.

Now let us move to the CAPM. In Chapter 3, we showed that the expected return of the asset is

$$E_S = r + (E_M - r)\beta_S \tag{5.9}$$

Substituting (5.6) and (5.8) into (5.9), we find that the expected return on the futures is

$$E_F = (E_M - r)\beta_F = (E_M - r)\beta_S \tag{5.10}$$

While on first appearance the relation (5.10) may seem perplexing, it makes a good deal sense intuitively. In buying the asset, we actually buy two things—the risk-free asset and a risk premium. We are entitled to the rate of return on the risk-free asset, r, because we have funds tied up in the asset, independent of its risk level. In addition, we are entitled to the risk premium associated with holding the asset, $(E_M - r)\beta_S$, because we have put our investment at risk. In buying the futures, we have accepted only the risk and, therefore, are entitled to receive only the risk premium, $(E_M - r)\beta_S$. With no funds tied up, we have no right to any risk-free return.

Relation to Net Cost of Carry

In Chapter 4, we discussed the net cost of carry relation and its implications. We showed that being long the asset and short a futures meant that we were implicitly long risk-free bonds. Equations (5.9) and (5.10) confirm this result. Being

[2] The rules of expectation operators are provided in Appendix A: Elementary Statistics at the end of the book.

TABLE 5.1 Perfect substitutes implied by the capital asset pricing model.

Position 1		Position 2
Buy asset/sell forward	=	Buy risk-free bonds (lend)
Buy risk-free bonds (lend)/buy forward	=	Buy asset
Buy asset/sell risk-free bonds (borrow)	=	Buy forward
Sell asset/buy forward	=	Sell risk-free bonds (borrow)
Sell risk-free bonds (borrow)/sell forward	=	Sell asset
Sell asset/buy risk-free bonds (lend)	=	Sell forward

long the asset means that we expect rate of return E_S and being short the futures means that we expect rate of return E_F. Thus the net portfolio return from being long the asset and short the futures equals the risk-free rate of interest,

$$E_S - E_F = r + (E_M - r)\beta_S - (E_M - r)\beta_S = r \tag{5.11}$$

The risk premium associated with buying the asset is exactly offset by the risk premium associated with selling the futures.

Just as in Chapter 4, we can pair up any two instruments to create the other. Suppose, for example, we buy risk-free bonds and buy a futures. The expected portfolio return is exactly equal to that of the underlying asset, that is,

$$r + E_F = r + (E_M - r)\beta_S = E_S \tag{5.12}$$

Table 5.1 summarizes all possible pairings, and is the counterpart to Table 4.5 in Chapter 4. The intuition is simple. Buying or selling a futures is the same as buying or selling a risk premium. Buying and selling an asset, on the other hand, means buying and selling a portfolio that consists of the risk-free asset and a risk premium. Note that, if the risk premium of the asset happens to equal zero, the expected rate of price change in the futures is zero.

Futures as Predictor of Expected Asset Price

The relation between expected return and risk of the asset and the futures also provides us with insight regarding the relation between the current futures price and the expected asset price when the futures expires at time T. To see this, consider committing to buy the asset at time T. The present value of the expected asset price is

$$S = E(\tilde{S}_T)e^{-E_S T}$$

where E_S is the asset's expected risk-adjusted rate of return.[3] On the other hand, consider committing to buy the asset at time T by buying a futures contract

[3] Recall that in Chapter 3 we used the CAPM to arrive at this value.

today at price F. Since F is paid at time T and is certain, the present value of this obligation is Fe^{-rT}. Since both quantities represent the same thing—the present value of one unit of the commodity at time T—they should be equal in value. Thus, the current futures price may be written

$$F = E(\tilde{S}_T)e^{-(E_S - r)T} \tag{5.13}$$

The structure of (5.13) says that the difference between the futures price and the expected asset price is nonzero. This means that the futures price is a biased predictor of the expected asset price. If the risk premium is positive, as is usually the case, the futures price is a downward biased predictor. The only instance in which the futures price is an unbiased predictor of the expected future asset price is where the risk premium of the asset equals 0.

Hedging Assets Using Futures Contracts

With the CAPM framework in hand, we can now turn to the exercise of managing the expected return and risk of a position in the asset underlying the futures contract. To do so, consider a portfolio that consists of one unit of the asset and futures contracts. Its expected rate of return is equal to

$$E_H = E_S - n_F E_F = (1 + n_F)E_S - n_F r \tag{5.14}$$

To find its total risk, recognize that the rate of return relation (5.5) implies that the futures return and the asset return are perfectly positively correlated, that is, $\rho_{SF} = +1$. This means that a portfolio that consists of one unit of the asset and n_F futures contracts has a standard deviation (i.e., total risk) equal to

$$\sigma_H = |\sigma_S + n_F\sigma_F| = |(1 + n_F)\sigma_S| \tag{5.15}$$

Managing expected return and risk of the portfolio therefore amounts to selecting a value for n_F.

Figure 5.1 summarizes some obvious choices of n_F. To make matters as a clear as possible, we use numerical values for the expected return and risk parameters. Specifically, we assume that the expected return and risk of the asset are 12% and 20%, respectively, and the risk-free rate of interest is 4%. This means that the expected return and risk of the futures contract are 8% and 20%. Where $n_F = 0$, the portfolio is unhedged. We hold only the asset, and the portfolio has an expected return of 12% and a risk of 20%. Selling futures against the long position in the asset reduces expected return and risk. At $n_F = -0.5$, we are implicitly selling one futures contract for every two units of the asset we hold. The risk level of this portfolio is below the risk level of the asset so we have "hedged." This particular hedge portfolio has an expected return of $0.5E_S + 0.5r$ or 8% and a risk of $0.5\sigma_S$ or 10%. Where we set $n_F = -1$, the hedge portfolio has an expected return of 4% and no risk. Since this is the lowest risk level possible, this particular hedge portfolio is called the "risk-minimizing hedge."

FIGURE 5.1 Relation between expected return and risk for portfolio consisting of one unit of asset and n_F futures contracts.

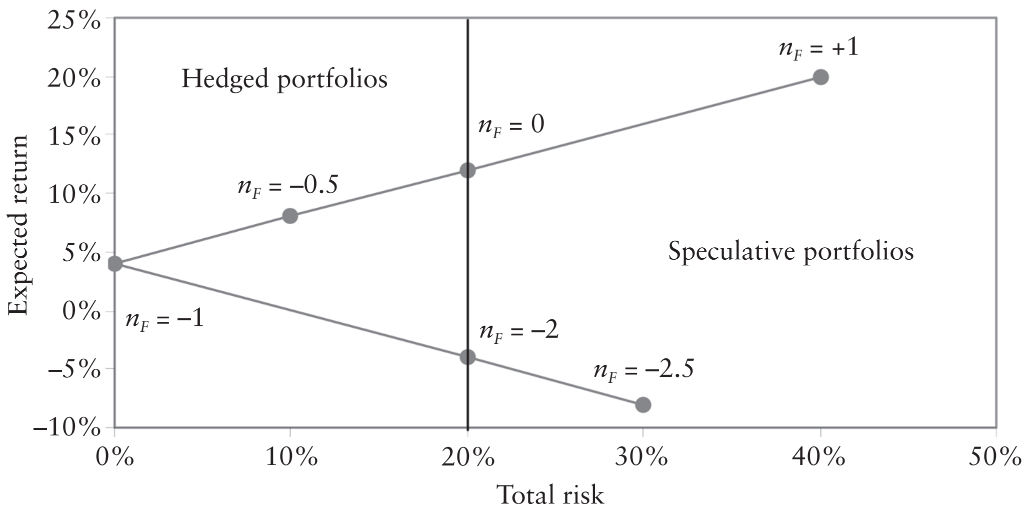

Note that, as we continue to sell more futures (i.e., $n_F < -1$), risk starts to increase, but, as long as $\sigma_H < \sigma_S$, we are hedging. Where $n_F < -2$, we are speculating in that the portfolio risk level exceeds the risk level of the asset held in isolation. The same applies where we buy futures (i.e., $n_F > 0$) rather than sell. Thus in Figure 5.1, where $-2 < n_F < 0$, we hold a hedged portfolio, and, where $n_F < -2$ and $n_F > 0$, we hold a speculative portfolio.

HEDGING PRICE RISK

In general, identifying the set of viable hedge opportunities is more complicated than Figure 5.1 suggests. The reason is *basis risk*. Basis risk refers to the fact that the futures price movements and asset price movements are not perfectly correlated. To understand why, it is useful to think of basis risk as being the sum of two components, that is,

$$\text{Basis risk} = \text{Time basis risk} + \text{Grade basis risk} \tag{5.16}$$

Time basis risk refers to uncertainty in the difference between the futures price and the underlying asset price. In the first section, the time basis risk was equal to zero because we assumed the net cost of carry relation holds at all points in time. In order to arrive at that relation we assumed that markets are frictionless and that the risk-free rate of interest, r, and the income rate on the asset, i, are constant through time. As a practical matter, arbitrageurs incur trading costs, and short-term interest and income rates may have a modest amount of uncertainty. This means that the futures price movements and asset price movements will not be perfectly correlated, except in the special case where the length of the hedge horizon exactly matches the time to expiration of the futures and the convergence of the futures and asset prices is assured.

FIGURE 5.2 Evolution of time and grade basis over the life of futures contract.

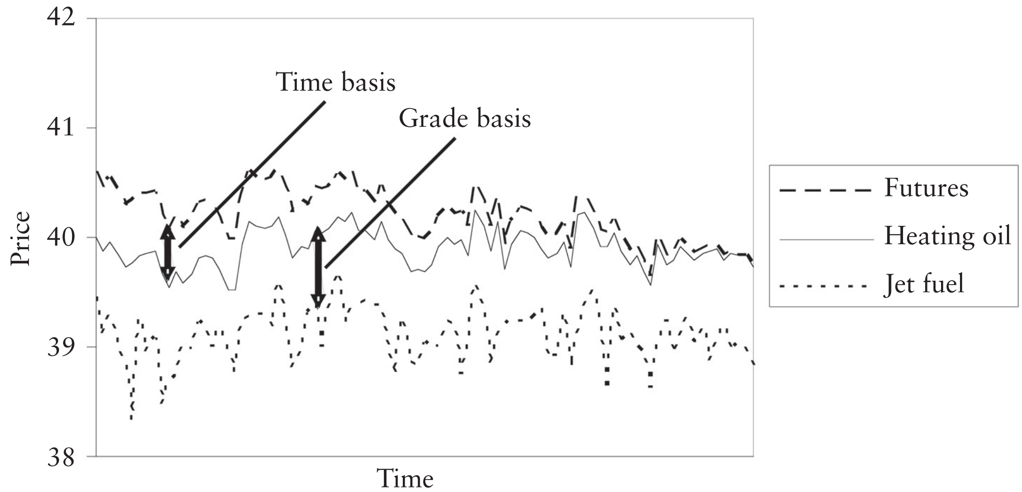

The second component is grade basis risk. Often we find situations in which futures contracts are not written on the asset whose price risk we want to manage. Many airlines, for example, want to hedge their jet fuel costs, however jet fuel futures contracts are not available. Fortunately, jet fuel and heating oil are very close substitutes, and heating oil futures can be used to cross-hedge. In this case, *grade basis risk* refers to the uncertainty in the difference between the price of heating oil and the price of jet fuel.[4] Figure 5.2 shows the evolution of time and grade basis over the life of the futures contract. The top line represents the heating oil futures prices, the middle line heating oil, and the bottom line jet fuel. The difference between the prices of the heating oil futures and heating oil is the time basis. As time passes, the time basis narrows. At expiration, the futures price equals the spot price of heating oil, and the time basis is zero. If the length of the hedging horizon is less than the life of the futures, a futures hedge must be unwound prior to expiration and time basis risk is incurred. The difference between the heating oil price and the jet fuel price (the lowest line) is the grade basis. It too varies through time. In this instance, however, convergence is not assured. Thus, in using heating oil futures to hedge the price of jet fuel, both time basis risk and grade basis risk are incurred. We now develop a framework for handling such a price risk management problem.

Minimize Price Risk

To make the development of a price risk-minimizing hedge as an intuitive as possible, let us use the example of an airline that wants to minimize the price risk of jet fuel that it needs at time *T*. Assume, for the sake of simplicity, that the

[4] The term "grade" arose in the agricultural futures market. The wheat futures contract traded on the Chicago Board of Trade, for example, allows the short futures to deliver different "grades" of wheat.

airline has no ability to store jet fuel—it buys jet fuel as it is needed at the market price. Assume also that we are considering a single refueling at time T. The jet fuel price at time T is denoted $\tilde{S}_T$. The current heating oil futures price is denoted F, and its price at time T is denoted $\tilde{F}_T$. Assuming the airline buys n_F futures contracts, its net cost of jet fuel at time T is

$$\tilde{C}_T = \tilde{S}_T + n_F(\tilde{F}_T - F) \tag{5.17}$$

Naturally, where $n_F = 0$, the airline pays the market price for fuel at time T.

To find the risk-minimizing hedge, use (5.17) to help write the variance of the hedged cost of fuel, that is,

$$\begin{aligned} \text{Var}(\tilde{C}_T) &= \text{Var}(\tilde{S}_T + n_F\tilde{F}_T) \\ &= \text{Var}(\tilde{S}_T) + n_F^2\text{Var}(\tilde{F}_T) + 2n_F\text{Cov}(\tilde{S}_T, \tilde{F}_T) \end{aligned} \tag{5.18}$$

where $\text{Var}(\tilde{S}_T)$ and $\text{Var}(\tilde{F}_T)$ are the variances of the asset and futures prices, respectively, and $\text{Cov}(\tilde{S}_T, \tilde{F}_T)$ is the covariance of the asset and futures prices. To find the number of futures contracts necessary to minimize $\text{Var}(\tilde{C}_T)$, n_F^*, we take the derivative of (5.18) with respect to n_F, and set it equal to zero, that is,

$$\frac{d\text{Var}(\tilde{C}_T)}{dn_F} = 2n_F^*\text{Var}(\tilde{F}_T) + 2\text{Cov}(\tilde{S}_T, \tilde{F}_T) = 0 \tag{5.19}$$

Rearranging, we find that the risk-minimizing hedge is

$$n_F^* = -\frac{\text{Cov}(\tilde{S}_T, \tilde{F}_T)}{\text{Var}(\tilde{F}_T)} = -\rho(\tilde{S}_T, \tilde{F}_T)\frac{\sqrt{\text{Var}(\tilde{S}_T)}}{\sqrt{\text{Var}(\tilde{F}_T)}} \tag{5.20}$$

where $\rho(\tilde{S}_T, \tilde{F}_T)$ or, simply ρ, is the correlation between the asset and futures prices.

The "optimal" hedge, n_F^*, as shown by the expression (5.20), is interesting in a number of respects. First, and foremost, n_F^* is negative since the variances are positive by definition and the correlation between the asset and futures prices is, presumably, positive. This means that, if the hedger is long the asset, he needs to sell futures, and vice versa. Second, if the futures is written on the specific asset being hedged, and the futures expires at the end of the hedge period, the end-of-period prices must be equal $\tilde{F}_T = \tilde{S}_T$. This implies that the variances of the asset and futures prices are equal, $\text{Var}(\tilde{S}_T) = \text{Var}(\tilde{F}_T)$, and that the correlation between the asset and futures prices is one, $\rho = +1$. The risk-minimizing futures hedge is therefore a one-to-one hedge against the asset, that is,

$$n_F^* = -\rho\frac{\sqrt{\text{Var}(\tilde{S}_T)}}{\sqrt{\text{Var}(\tilde{F}_T)}} = -\frac{\sqrt{\text{Var}(\tilde{S}_T)}}{\sqrt{\text{Var}(\tilde{F}_T)}} = -1$$

Third, the effectiveness of the hedge depends upon the correlation between the asset and futures prices. If $\rho = +1$, the hedge is perfect, and the optimal hedge is to sell one futures contract. If $\rho = -1$, the hedge is also perfect, and the optimal hedge is to buy one futures contract. If $-1 < \rho < +1$, the hedge will not be fully effective, with the effectiveness decreasing as the correlation approaches 0. At $\rho = 0$, the asset and futures prices are independent, so there is no point in taking a futures position.

Estimating Variance/Covariance

Before applying the risk-minimizing hedge framework, we need to discuss how to estimate the variance and covariance expressions in (5.20). In deriving the risk-minimizing hedge, we formulated the problem as a one-period hedge, from time 0 to time *T*. The length of the hedge period is arbitrary. For the sake of illustration, assume its *T* days. Now let us consider the variance of the asset price, $\text{Var}(\tilde{S}_T)$. Over the hedge horizon, we will observe a sequence of asset prices $S, \tilde{S}_1, \tilde{S}_2, \ldots, \tilde{S}_T$. To see how $\text{Var}(\tilde{S}_T)$ can be expressed in terms of the price sequence, note that

$$\text{Var}(\tilde{S}_T) = \text{Var}(\tilde{S}_T - S_0) = \text{Var}\left(\sum_{t=1}^{T} \tilde{S}_t - \tilde{S}_{t-1}\right) = \text{Var}\sum_{t=1}^{T} \Delta\tilde{S}_t \tag{5.21}$$

where $\Delta\tilde{S}_t = \tilde{S}_t - \tilde{S}_{t-1}$. Assuming that price changes are independent and identically distributed (i.i.d.), the variance of the end-of-period asset price is simply *T* times the daily variance of the asset price change, that is,

$$\text{Var}(\tilde{S}_T) = T\text{Var}(\Delta\tilde{S}) \tag{5.22}$$

By the same logic, the variance the end-of-period futures price is

$$\text{Var}(\tilde{F}_T) = T\text{Var}(\Delta\tilde{F}) \tag{5.23}$$

and the covariance of the end-of-period asset price and futures price is

$$\text{Cov}(\tilde{S}_T, \tilde{F}_T) = T\text{Cov}(\Delta\tilde{S}, \Delta\tilde{F}) \tag{5.24}$$

Thus, the risk-minimizing hedge over the interval from 0 to *T* (5.20) can be rewritten in terms of daily price changes

$$\begin{aligned} n_F^* &= -\frac{\text{Cov}(\tilde{S}_T, \tilde{F}_T)}{\text{Var}(\tilde{F}_T)} = -\frac{T\text{Cov}(\Delta\tilde{S}, \Delta\tilde{F})}{T\text{Var}(\Delta\tilde{F})} \\ &= -\frac{\text{Cov}(\Delta\tilde{S}, \Delta\tilde{F})}{\text{Var}(\Delta\tilde{F})} = -\rho_{\Delta S, \Delta F}\frac{\sigma_{\Delta S}}{\sigma_{\Delta F}} \end{aligned} \tag{5.25}$$

In the hedge illustrations developed through this chapter, we take advantage of this property.

Setting Risk-Minimizing Hedge

We now show how to apply the risk-minimizing hedge framework. Assume that the airline needs 150,000 gallons of jet fuel in exactly 30 days and that it wants to minimize the variance of the cost of acquiring the fuel. Assume also that the airline has no ability to store jet fuel. Setting a risk-minimizing hedge has four steps.

Step 1: Identify Appropriate Futures Contract Choosing the appropriate futures contract involves at least two factors. First, the higher is the correlation between the futures price and the fuel price, the more effective is the hedge. Ideally this means using a jet fuel futures contract to hedge, if one is available. In this way, only time basis risk is incurred. As noted earlier, however, futures contracts on jet fuel are not traded. The closest substitute is heating oil futures. Second, given heating oil futures listed on the New York Mercantile Exchange (NYMEX)[5] have 18 different contract maturities, how do we choose among available contracts? The tradeoff here is contract liquidity versus the cost of "rolling" the futures position. In general, nearby contracts are the most liquid and, hence, have the lowest trading costs. Unfortunately, however, the nearest available contract may expire before the hedge horizon is complete, in which case we must roll into the next available maturity (i.e., the nearby futures position is closed and a second nearby futures position is entered). Given that the hedge horizon is only 30 days in our illustration, using the heating oil futures contract that expires just after the hedge horizon is compete probably makes the most sense.

Step 2: Collect Historical Prices With the futures contract selected, we must now estimate the standard deviations and correlation of the daily price changes on the right-hand side of (5.25). Note that, within the hedge framework, these values are expected *future* standard deviations and correlation. Since we have no means of observing these parameters, we usually rely on historical time-series data to develop estimates. For the problem at hand, we will have to collect historical time series data for the jet fuel and the heating oil futures contract. These data are provided in the Excel file, **Jet fuel.xls**.

Step 3: Estimate Standard Deviation and Correlation Parameters With the data in hand, we now compute the standard deviation of the jet fuel price change, the standard deviation of the heating oil futures contract price change, and the correlation between the price changes of jet fuel and the heating oil futures. The estimator of the standard deviation of the historical asset price change series is

$$\hat{\sigma}_{\Delta S} = \sqrt{\frac{1}{T-1}\sum_{t=1}^{T}(\Delta S_t - \overline{\Delta S})^2} \qquad (5.26)$$

where the symbol "^" indicates a specific estimate based on a sample of prices and T is the number of historical prices in the time series ($t = 1, \ldots, T$).[6] The

[5] The NYMEX is the dominant exchange in the U.S. listing futures contracts on petroleum and petroleum products.

[6] These formulas are taken from the review of elementary statistics provided in Appendix A of this book.

standard deviation of the historical future price change series is similar. The estimator of the correlation between two historical price change series of the asset and the futures is

$$\hat{\rho}_{\Delta S,\,\Delta F} = \frac{\sqrt{\sum_{t=1}^{T}(\Delta S_t - \overline{\Delta S})(\Delta F_t - \overline{\Delta F})}}{\sqrt{\sum_{t=1}^{T}(\Delta S_t - \overline{\Delta S})^2 \sum_{t=1}^{T}(\Delta F_t - \overline{\Delta F})^2}} \tag{5.27}$$

The estimates are: $\hat{\sigma}_{\Delta S} = 0.0422$, $\hat{\sigma}_{\Delta F} = 0.0357$, and $\hat{\rho}_{\Delta S,\,\Delta F} = 0.9320$.

Step 4: Compute the Risk-Minimizing Hedge The fourth and final step is to compute the optimal number of futures, which is done using (5.25), that is,

$$n_F^* = -\frac{0.9320(0.0422)}{0.0357} = -1.0997$$

The negative sign implies that we sell futures to create the risk-minimizing hedge. The optimal number of futures to sell is –1.0997 gallons for each gallon of jet fuel. Figure 5.3 shows the effect that changing the number of futures has on the standard deviation of cost per gallon.

Once the optimal hedge ratio is determined, finding the number of futures contracts to use is a matter of multiplying the hedge ratio by the quantity

FIGURE 5.3 Relation between the number of futures contracts held (+ long; – short) and the risk (standard deviation) of hedged jet fuel cost. (Parameters: $\hat{\sigma}_{\Delta S} = 0.0422$, $\hat{\sigma}_{\Delta F} = 0.0357$, and $\hat{\rho}_{\Delta S,\,\Delta F}\, \rho_{\Delta S,\Delta F} = 0.9320$.)

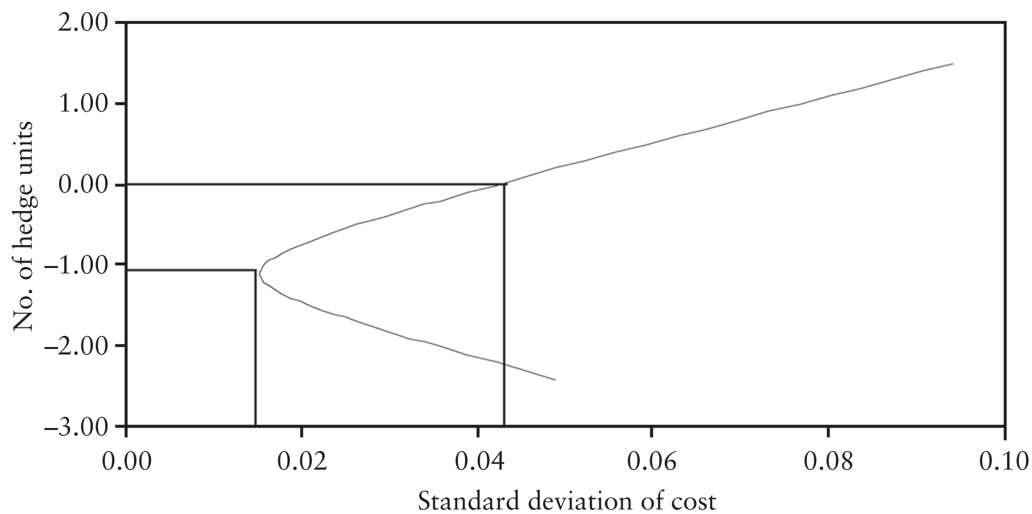

demanded and dividing by the futures contract denomination. A single heating oil futures has a 42,000 gallon denomination, so the number of futures to sell is

$$1.997\left(\frac{150{,}000}{42{,}000}\right) = 3.928$$

In practice, we would also "tail the hedge" by multiplying 3.928 by the discount factor $e^{-r(30/365)}$. One each subsequent day, the number of contracts would be increased a a factor of $e^{r(1/365)}$. For expositional convenience, we ignore this practice of tailing the hedge through the remainder of the chapter.

Relation to OLS Regression

Ordinary least squares (OLS) regression offers a convenient direct means of estimating the risk-minimizing hedge. Once the necessary time-series price data (t = 1, . . ., T) are collected, run the regression,

$$\Delta\tilde{S}_t = \alpha_0 + \alpha_1 \Delta\tilde{F}_T + \tilde{\varepsilon}_t \tag{5.28}$$

As it turns out, the value of the regression coefficient α_1 under OLS regression is defined as[7]

$$\alpha_1 = \frac{\text{Cov}(\Delta\tilde{S}, \Delta\tilde{F})}{\text{Var}(\Delta\tilde{S})} \tag{5.29}$$

Thus the risk-minimizing number of futures contracts may be written as a function of the estimated slope coefficient, that is,

$$n_F^* = -\alpha_1 \tag{5.30}$$

Aside from eliminating the need to estimate directly the individual standard deviations and correlation in (5.25), the OLS regression's adjusted R-squared provides a measure of hedging effectiveness. The risk of the unhedged asset price risk is $\text{Var}(\tilde{S}_T)$. Of this amount, selling α_1 futures explains $\alpha_1^2\text{Var}(\tilde{F}_T)$. The adjusted R-squared therefore tells us the percent of the total unhedged asset price change risk that is hedgable risk,

$$\begin{aligned}\bar{R}^2 &= \frac{\alpha_1^2\text{Var}(\tilde{F}_T)}{\text{Var}(\tilde{S}_T)} = \frac{\text{Hedgable risk}}{\text{Unhedged risk}} \\ &= 1 - \frac{\text{Var}(\tilde{\varepsilon}_T)}{\text{Var}(\tilde{S}_T)} = 1 - \frac{\text{Unhedgable risk}}{\text{Unhedged risk}}\end{aligned} \tag{5.31}$$

[7] See OLS regression review is Appendix B of this book.

An adjusted R-squared of 100% means that the asset price risk is fully hedgable, while an adjusted R-squared of 0% says that there is no point in hedging.

To verify that the ordinary least squares regression approach produces the same risk-minimizing hedge, apply the OPTVAL Library function,

OV_STAT_OLS_SIMPLE(y, x, *intercept*, *out*)

where y is the vector of jet fuel prices, x is the vector of heating oil futures prices, *intercept* is an indicator variable whose value is "Y" is the regression includes an intercept term and is "N" is the intercept term is being suppressed, and *out* is an indicator variable instructing the function to display the results horizontally ("H" or "h") or vertically ("V" or "v"). This function returns a horizontal array of output of length 5. The array contains the estimate of the intercept terms, its standard error, the estimate of the slope coefficient, its standard error, and the adjusted R-squared. When calling this function, we must highlight five contiguous cells, enter the relevant data, and then press the Shift, Ctrl, and Enter keys simultaneously. The panel below demonstrates.

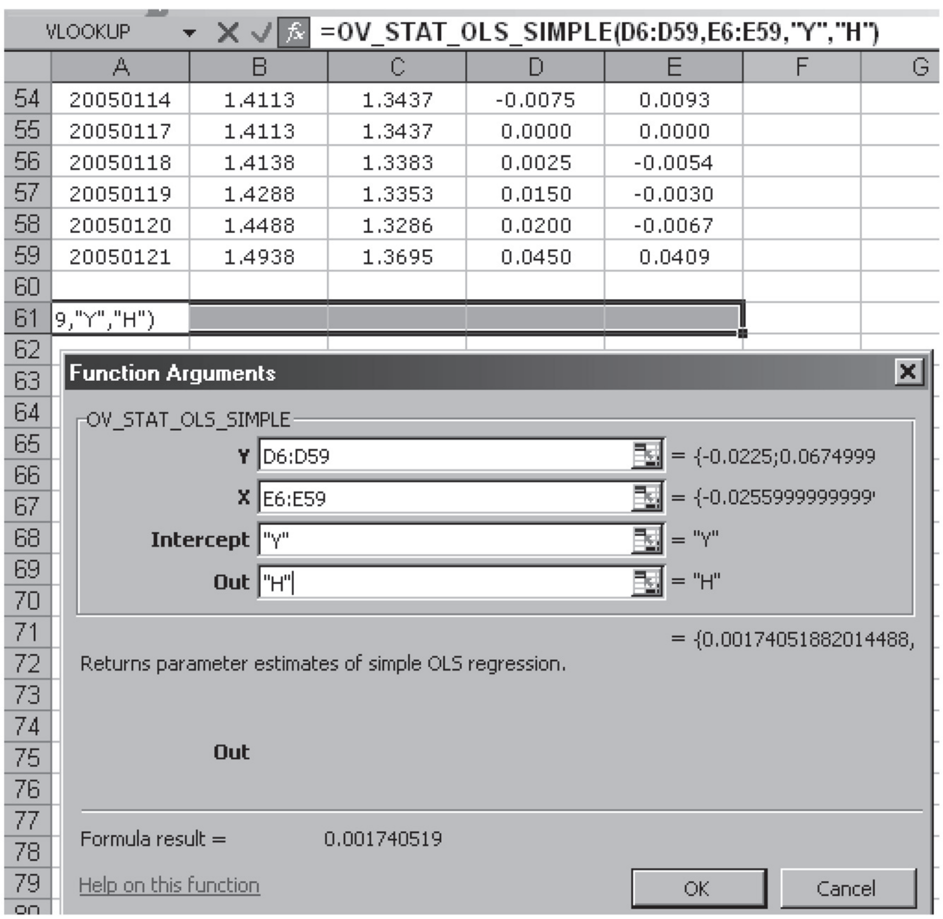

The estimated slope coefficient is 1.0997, exactly as before.

Criterion	$\hat{\alpha}_0$	$s(\hat{\alpha}_0)$	$\hat{\alpha}_1$	$s(\hat{\alpha}_1)$	$\bar{R}^2$
Minimize price risk	0.0017	0.0021	1.0997	0.0593	0.8686

The focus on price changes (and returns) rather than prices in risk management becomes the norm from this point forward. Where the underlying asset is a physical asset or commodity like grain, we assume that price changes are independent and identically distributed (i.i.d.) through time. The reason is that most commodity prices tend to be mean-reverting. The reason is simple. If the price of a commodity becomes too low, producers of the commodity will slow or stop production, inventories will become depleted, and prices will rise. If the price of a commodity becomes too high, consumers will cut back on demand, and prices will fall. The prices of financial assets, however, are different. Consider a common stock. The company engages in a particular type of business activity and generates cash flow. This cash flow is used to expand operations, and the expanded operations generate proportionately more cash. For such an asset, price, it is more reasonable to assume that price is expected to grow at a constant rate and to have a constant variance rate. To model such behavior, it is most common to assume that the difference in the natural logarithm of asset prices or continuous returns (i.e., $\ln(S_t) - \ln(S_{t-1}) = \ln(S_t/S_{t-1}) = R_t$) are independent and identically distributed (i.i.d.) through time.

HEDGING REVENUE RISK

The apparatus for managing price risk is the same independent of whether we are managing the risk of costs or income. A corn farmer, for example, may want to hedge the price at which he will sell his crop when he harvests in the fall. To identify the risk-minimizing hedge, he can run a regression of the price per bushel of the grade of corn that he has planted on the price per bushel of a corn futures contract.[8] He would then multiply the estimated slope coefficient (i.e., the risk-minimizing hedge per bushel) by his planned harvest size and divide by the futures contracts size (i.e., 5,000 bushels) to determine the number of contracts to sell.

This oversimplifies the farmer's problem, however. When he seeds his fields in the spring, both the price of corn, $\tilde{S}_T$, and the yield per acre, $\tilde{n}_T$, at the time of harvest in the fall are unknown. What is more germane to the farmer is revenue. In all likelihood, he is more interested in minimizing the revenue risk (i.e., the product of price and quantity),

$$\mathrm{Var}(\tilde{R}_T) = \mathrm{Var}(\tilde{n}_T\tilde{S}_T + n_F\tilde{F}_T) \tag{5.32}$$

rather than price risk alone. It is interesting to note that the relation price between and yield provide, to some degree, a *natural* hedge. If weather conditions are poor during the summer months, the harvest size will be small and the price per bushel will likely to be high. On the other hand, the fall brings a

[8] The corn futures contract traded at the Chicago Board of Trade calls for the delivery of No. 2 yellow corn at par, No. 1 yellow corn at 1½ cents per bushel over the contract price, or No. 3 yellow at 1½ cents per bushel under the contract price. Assuming the farmer has planted yet a different grade, he incurs both time and grade basis risk.

bumper crop, prices are likely to be low. The negative correlation between price and quantity tends to reduce the level of revenue risk, holding other factors constant.

Whether the correlation is negative or positive is irrelevant in estimating the revenue risk-minimizing hedge. To find the revenue risk-minimizing hedge, we simply replace the dependent variable in regression (5.28). In place of using the price change of corn, we use the change in revenue per acre.

ILLUSTRATION 5.1 Hedging price risk versus revenue risk.

Consider the case of a farmer who has just planted his 10,000 acres of land with wheat. Compare the number of futures contracts to sell if he decides to minimize revenue risk rather than price risk. The Excel file, ***Wheat.xls****, contains historical data over the past 30 years. For the price risk-minimizing hedge, assume the farmer anticipates harvesting 60 bushels per acre.*

The first step is to summarize the data, and compute revenue per acre. The figures are shown in the table below. The average wheat price at harvest over the past 30 years was \$3.00 per bushel, and the average yield per acre was 60 bushels. The "Revenue per acre" column with the subheading "Constant yield" is simply the harvest price times 60 bushels per acre (e.g., 2.469 × 60 = 148.15 per acre), and the "Revenue per acre" column with the subheading "Varying" is the harvest price times yield per acre in that year (e.g., 2.469 × 68.39 = 168.86 per acre).

				Revenue per Acre		Change in Revenue per Acre		
Month	Spot Price	Futures Price	Yield	Constant Yield	Varying Yield	Constant Yield	Varying Yield	Futures Price Change
1	2.469	2.448	68.39	148.15	168.86			
2	2.664	2.638	64.76	159.82	172.50	11.67	3.64	0.191
3	2.176	2.123	70.32	130.56	153.02	–29.26	–19.48	–0.515
4	2.481	2.501	69.08	148.88	171.41	18.32	18.40	0.378
5	2.737	2.686	65.20	164.24	178.47	15.36	7.06	0.185
…	…	…	…	…	…	…	…	…
26	3.493	3.667	57.66	209.61	201.42	0.93	41.24	–0.260
27	3.700	3.620	50.52	222.00	186.91	12.40	–14.51	–0.047
28	4.065	4.054	45.29	243.88	184.11	21.88	–2.80	0.434
29	3.833	4.054	50.68	230.00	194.26	–13.88	10.15	0.000
30	3.298	3.722	53.68	197.88	177.02	–32.12	–17.23	–0.332
Mean	3.000	3.017	60.00	180.00	176.61	1.71	0.28	
StDev	0.480	0.557	7.887	28.77	11.00	17.70	15.20	

To find the price risk-minimizing hedge, we regress the "Change in revenue per acre—Constant yield" column on the "Futures price change" column, and, to find the revenue risk-minimizing hedge we regress the "Change in revenue per acre—Varying yield" column on the "Futures price change" column. The estimated slope coefficients, $\hat{\alpha}_1$, in the regressions are the number of bushels of wheat that need to be sold using the futures contract. The regression results are as follows:

Criterion	$\hat{\alpha}_0$	$s(\hat{\alpha}_0)$	$\hat{\alpha}_1$	$s(\hat{\alpha}_1)$	$\bar{R}^2$
Minimize price risk	–0.2765	1.7110	45.3321	5.1264	0.7433
Minimize revenue risk	0.1242	2.8909	3.5806	8.6616	0.0063

The regression results reveal that the two hedges are quite different from each other. If the farmer chooses to minimize price risk, he needs to sell 45.3321 bushels in futures per acre of land. With 10,000 acres, this means a total of 453,321 bushels. The wheat futures contracts traded on the Chicago Board of Trade have a denomination of 5,000 bushels, so a total of 90.664 contracts should be sold. On the other hand, if the farmer chooses to minimize revenue risk, he needs to sell only 3.5806 bushels per acre, or 7.161 futures contracts. Fewer contracts are required in the revenue risk-minimizing hedge because price and yield per care are inversely related. This negative correlation manifests itself in risk exposure. In the above table, the standard deviation of the revenue change with the fixed 60 bushels per acre (i.e., price risk) is 17.70, while the standard deviation of revenue change with varying yield (i.e., revenue risk) is 15.20. Because price and quantity are inversely related, the amount of risk that needs to be managed is less.

HEDGING MARGIN RISK

Another type of risk that may be faced by a processor or producer is gross processing margin risk. *Gross processing margin* refers to the difference between total revenue from production and the total costs of production, that is,

$$\tilde{M}_T = \tilde{n}_O \tilde{S}_{O,T} - n_I \tilde{S}_{I,T} - \text{Fixed costs} \tag{5.33}$$

where $n_{O,T}$ is the quantity demanded at time T when output price is $S_{O,T}$ per unit, n_I is the number of input units required for production, and $\tilde{S}_{I,T}$ is input cost per unit.[9] Consider an oil refiner, for example. In the normal course of production, he buys crude oil, distills it, and sells heating oil and unleaded gasoline. If he is planning for production that will occur at time T, he faces both revenue and price risk. The revenue risk arises because the refiner knows neither the market price per gallon of product (e.g., unleaded gasoline) at time T nor the number of gallons that will be demanded. The price risk arises because the price per barrel of crude oil at time T will depend on supply and demand conditions at that time. Thus the refiner's risk management problem may be to minimize the variance of his margin risk,

$$\text{Var}(\tilde{n}_O \tilde{S}_{O,T} - n_I \tilde{S}_{I,T} + n_F \tilde{F}_T) \tag{5.34}$$

Like in the previous examples, this can be accomplished by regressing the change in the gross processing margin on the futures price change. The slope coefficient estimate realized from the regression in the risk-minimizing number of futures. It is

[9] We assume that production takes place instantaneously (i.e., products are produced as quickly as the inputs are acquired), and that unsold production cannot be carried over from one period to the next. Naturally, both of these assumptions can be relaxed.

important to recognize that, like a revenue hedge, a margin hedge implicitly accounts for the fact that input costs and output prices may be strongly correlated. Such a natural hedge reduces the number of necessary futures contracts.

ILLUSTRATION 5.2 Hedging margin risk.

Consider the case of a gold watch manufacturing firm. Over the past 67 months, they have produced and sold an average of 6,274 gold watches per month at an average sales price of $3,727. Month-by-month sales statistics are included in ***Watch manufacturer.xls****. The key input cost of each watch is gold, and its price is uncertain from month to month. All the firm knows is that it takes four Troy ounces of gold to manufacture each watch. Their fixed monthly costs are $5,000,000. Find the optimal number of futures contract to enter to minimize the variance of the firm's end-of-month profit margin. Also, compute the minimum revenue risk and minimum cost risk hedges. The denomination of the gold futures contract is 100 Troy ounces.*

The data file contains a history of sales prices and quantity sold together with prices of gold and gold futures over a 67-month period. The format is as follows:

		Gold Watch Production				
Month	Gold	Price	Quantity	Revenue	Gold Cost	Margin
19990101	287.75	3,444.45	6,534	22,504,718	–7,520,201	9,984,517
19990201	287.65	3,467.84	6,511	22,578,920	–7,491,500	10,087,419
19990301	285.85	3,455.90	6,546	22,623,962	–7,485,246	10,138,716
19990401	280.45	3,381.07	6,600	22,316,724	–7,404,423	9,912,301
…	…	…	…	…	…	…

		Changes in			
Month	Gold	Revenue	Gold Cost	Margin	Futures
19990101	287.75				
19990201	287.65	74,202	28,701	102,903	–2.21
19990301	285.85	45,042	6,255	51,297	–3.45
19990401	280.45	–307,238	80,823	–226,416	–4.93
…	…	…	…	…	…

The column headings are largely self-explanatory. Revenue equals the watch price times the quantity of watches sold. The cost of the gold used in each watch is the spot price of gold per ounce times four ounces times the quantity of watches produced. The margin equals the revenue less the variable costs less the $5,000,000 in fixed costs. The final four columns are the monthly changes in each of the variables.

Based on the information in the file, you regress (a) revenue change on the futures price change, (b) cost change on the futures price change, and (c) margin change on the futures price change. The results are summarized in the table below. The estimated slope coefficients, $\hat{\alpha}_1$, are the number of ounces of futures contracts that you should sell to hedge the effects of the gold price. Recall that the gold futures contract denomination is 100 ounces, so these figures need to be divided by 100.

Criterion	$\hat{\alpha}_0$	$s(\hat{\alpha}_0)$	$\hat{\alpha}_1$	$s(\hat{\alpha}_1)$	$\bar{R}^2$
Minimize revenue risk	19,314	28,729	24,462	2,028	0.6944
Minimize price risk	–5,705	9,484	–8,617	670	0.7213
Minimize margin risk	13,609	22,173	15,845	1,565	0.6155

To minimize the effect that gold price uncertainty has on revenue, you should sell 244.62 gold futures contracts. To minimize the effect that gold price uncertainty has on input costs, you should buy 86.17 contracts. Finally, to minimize the effect that gold price uncertainty has on profit margin, you should sell 158.45 gold futures. Note that the revenue risk hedge less the price risk hedge equals the margin hedge, 24,462 – 8,617 = 15,845. Apparently, the firm is able to pass along some the change in gold input cost by changing the price of its watches.

HEDGING PORTFOLIO VALUE

Up to this point in the chapter, we have looked at expected return/risk management of commodity price risk exposures embedded within operating costs, revenue, and gross margin. The next series of applications focus on managing the risk of financial assets. Suppose that we hold a portfolio of AAA-rated corporate bonds, for example, and know that there will be a major announcement by the Federal Reserve next week. Given the impending announcement, we would like to hedge our long-term interest rate risk exposure. One possible action is to liquidate our bond position. This action may be expensive, however, because the bond markets are not particularly liquid and trading costs are high. Another is to hedge the portfolio value using long-term interest rate futures contracts. Such markets are very liquid and trading costs are low. This section examines the expected return/risk management of a portfolio of securities where security value has only one source of underlying financial uncertainty (e.g., long-term interest rate risk, stock market risk, or currency risk).

To begin, we assume that the fund manager's objective function is to minimize the variance of the value of his portfolio over a single period ending at time T. The expression that we use for portfolio value risk is

$$\text{Var}(\tilde{V}_T + n_F \tilde{F}_T) \tag{5.35}$$

where V_T is the sum of the market values of all securities in the portfolio at time T, and F_T is the price of the futures contract most closely tied to the portfolio's underlying source of risk (e.g., if V_T is a well-diversified portfolio of stocks, F_T would be a stock index futures contract).

Like in the previous risk-management problems of this chapter, we will focus initially on determining the risk-minimizing hedge. This approach to solving the problem needs to be modified slightly. The reason is that financial assets, unlike commodities, tend to grow in value through time. Consider a stock index portfolio, for example. Contributing to the growth in the value of this portfolio is the fact that not only do the individual stocks in the portfolio have prices that are expected to grow through time, but also any stocks that pay dividends will have those dividends reinvested in more shares of the stock portfolio. To manage this rate growth over the hedge horizon, we focus the natural logarithm of the portfolio value instead of the value itself. The sequence of the values over the hedge horizon is $\ln V, \ln \tilde{V}_1, \ln \tilde{V}_2, \ldots, \ln \tilde{V}_T$. Note that the end-of-hedge-period $\mathrm{Var}(\ln \tilde{V}_T)$ can be expressed in terms of the day to day values through time, that is,

$$\begin{aligned}\mathrm{Var}(\ln \tilde{V}_T) &= \mathrm{Var}(\ln \tilde{V}_T - \ln V_0) \\ &= \mathrm{Var}\left(\sum_{t=1}^{T} \ln \tilde{V}_t - \ln \tilde{V}_{t-1}\right) = \mathrm{Var}\left(\sum_{t=1}^{T} \tilde{R}_{V,t}\right)\end{aligned} \tag{5.36}$$

where $\tilde{R}_{V,t} = \ln(\tilde{V}_t / \tilde{V}_{t-1})$ is the continuously compounded return on the portfolio. Assuming that returns are independent and identically distributed (i.i.d.), the variance of the end-of-period portfolio value is simply T times the daily variance of the asset price change, that is,

$$\mathrm{Var}(\ln \tilde{V}_T) = T\mathrm{Var}(\tilde{R}_V) \tag{5.37}$$

By the same logic, the variance the end-of-period futures price is

$$\mathrm{Var}(\ln \tilde{F}_T) = T\mathrm{Var}(\tilde{R}_F) \tag{5.38}$$

and the covariance of the end-of-period asset price and futures price is

$$\mathrm{Cov}(\ln \tilde{V}_T, \ln \tilde{F}_T) = T\mathrm{Cov}(\tilde{R}_V, \tilde{R}_F) \tag{5.39}$$

Thus the risk-minimizing hedge over the interval from 0 to T (5.20) can be rewritten in terms of daily price changes

$$\begin{aligned}n_F^* &= -\frac{\mathrm{Cov}(\ln \tilde{V}_T, \ln \tilde{F}_T)}{\mathrm{Var}(\ln \tilde{F}_T)} = -\frac{T\mathrm{Cov}(\tilde{R}_V, \tilde{R}_F)}{T\mathrm{Var}(\tilde{R}_F)} \\ &= -\frac{\mathrm{Cov}(\tilde{R}_V, \tilde{R}_F)}{\mathrm{Var}(\tilde{R}_F)} = -\rho_{V,F}\left(\frac{\sigma_V}{\sigma_F}\right)\end{aligned} \tag{5.40}$$

where σ_V and σ_F are the standard deviations of the continuously compounded returns of the portfolio and the futures, and $\rho_{V,F}$ is the correlation between the rates of return of the portfolio and the futures.

Just as was the case in the earlier risk-minimizing hedge problems, the number of futures contracts to sell can be determined by OLS regression. Consider the relation between the portfolio value and the futures at the end of the hedge period standing today, that is,

$$\ln \tilde{V}_T = \alpha_0 + \alpha_1 \ln \tilde{F}_T + \tilde{\varepsilon}_T \tag{5.41}$$

The slope coefficient in this relation, α_1, is a price elasticity. It gives us the percentage change in the value of the portfolio for a given percentage change in the futures price. To hedge, however, we need to know dV/dF, that is, the change in the dollar value of the portfolio associated with a change in the futures price. In this way, we can sell exactly dV/dF futures contracts so that, if something unexpected happens and the value of the portfolio (and futures price) changes, the overall portfolio value change is zero, that is,

$$dV - n_F dF = dV - \left(\frac{dV}{dF}\right)dF = 0 \tag{5.42}$$

But the regression relation (5.41) provides only

$$\frac{d\ln V}{d\ln F} = \alpha_1$$

How can we get dV/dF? The answer in the chain rule:

$$\frac{dV}{dF} = \frac{dV}{d\ln V} \times \frac{d\ln V}{d\ln F} \times \frac{d\ln F}{dF} = V \times \alpha_1 \times \frac{1}{F} = \alpha_1\left(\frac{V}{F}\right) \tag{5.43}$$

We simply scale the regression coefficient α_1 by the ratio of the portfolio value to the value of a single futures contract.

To estimate α_1, we rely on the differenced form of (5.41), that is, the regression,

$$R_{V,t} = \alpha_0' + \alpha_1 R_{F,t} + \varepsilon_t' \tag{5.44}$$

where $\tilde{R}_{V,t} = \ln(\tilde{V}_t/\tilde{V}_{t-1})$ and $\tilde{R}_{F,t} = \ln(\tilde{F}_t/\tilde{F}_{t-1})$. We do this because the returns of the portfolio and the futures, $\tilde{R}_{V,t}$ and $\tilde{R}_{F,t}$, as well as the error term in the regression, ε_t', are assumed to be i.i.d. While the intercept term in (5.44) is not the same as the intercept term in (5.41), the slope coefficient (and, hence, the hedge ratio) is identical.

ILLUSTRATION 5.3 Hedging value.

Consider the case of a life insurance company that holds a large portfolio of AAA-rated corporate bonds. Its daily values for the period January 1, 2004 through February 16, 2005 are reported in ***Life insurance.xls****. Based on these values, the natural logarithm of value and portfolio return are computed. Also included in the file are the continuously-compounded returns of the Chicago Board of Trade's (CBT's) Treasury bond futures contract. The current (2/16/05) T-bond futures price is 1.1525 per dollar of face value, and the T-bond futures contract denomination is $100,000. Find the risk-minimizing hedge for the bond portfolio.*

The data included in the file are as follows:

Date	AAA Portfolio Value	Natural Log of Value	Portfolio Return	Futures Return
20040101	29,004,133	17.183		
20040102	28,677,998	17.172	–0.0113	–0.01295
20040105	28,679,125	17.172	0.0000	–0.00029
20040106	28,931,665	17.180	0.0088	0.01066
20040107	29,006,580	17.183	0.0026	0.00343
…	…	…	…	…
20050210	32,858,283	17.308	–0.0116	–0.01018
20050211	32,802,473	17.306	–0.0017	–0.00189
20050214	32,955,513	17.311	0.0047	0.00350
20050215	32,821,210	17.307	–0.0041	–0.00323
20050216	32,671,455	17.302	–0.0046	–0.00514

Based on the data, we regress the portfolio return on the futures return. This can be handled using the OPTVAL function, OV_STAT_OLS_SIMPLE. The results are as follows:

Criterion	$\hat{\alpha}_0$	$s(\hat{\alpha}_0)$	$\hat{\alpha}_1$	$s(\hat{\alpha}_1)$	$\overline{R}^2$
Minimize value risk	0.0001	0.0000	0.8935	0.0075	0.9797

The slope coefficient estimate, 0.8935, implies that, for a one percentage change in the futures price, the portfolio value will change by 0.8935%. To determine the hedge that minimizes the dollar value change of the portfolio, we must account for the current value of the portfolio as well as the price of the futures. Earlier in this illustration, we reported that the current portfolio value is \$32,671,455, and that the current futures price is 1.1525 times its \$100,000 denomination. Thus, the value risk-minimizing hedge is to sell 253.30 futures contracts:

$$n_F = -0.8935\left(\frac{32{,}671{,}455}{1.1525 \times 100{,}000}\right) = -253.30$$

With an adjusted R-squared of nearly 98%, you have good reason to believe that the futures hedge will be very effective.

HEDGING MULTIPLE SOURCES OF RISK

Aside from its convenience, the regression approach to setting a risk-minimizing hedge is easily generalized to handle asset portfolios whose values are influenced by a number of risk factors. Suppose we are managing a fund that invests primarily in stocks from the oil refining industry. Since the portfolio is not well diversified due to its concentration in oil stocks, its value is vulnerable not only to unexpected stock market movements but also to unexpected changes in the price of oil. Suppose that, given the political situation in Iraq, we come to the conclusion that there is a substantial risk that the price of crude oil will spike upward in the near future. This places us in a conundrum. While an increase in the crude oil price will likely cause the stock market level to fall, it may well have a positive influence on the prices of the oil stocks in our portfolio. Thus selling the stocks and buying risk-free bonds is not an appropriate strategy since it would eliminate both the stock market and crude oil price risk exposures. Our objective is to negate the stock market risk of our portfolio without negating the crude oil price risk.

A straightforward approach to handling this risk management problem is to use the multiple regression model,

$$R_V = \alpha_0 + \alpha_1 R_{F,1} + \alpha_2 R_{F,2} + \ldots + \alpha_n R_{F,n} + \tilde{\varepsilon} \tag{5.45}$$

where all futures contracts whose returns are thought to influence the value of our portfolio are used. With respect to the illustration at hand, we might include only two risk factors: the S&P 500 futures contract to proxy for stock market risk, and an oil futures contract to proxy for the effects of oil price risk. Once the regression is estimated, we can hedge any of the risk exposures using the estimated slope coefficients.

ILLUSTRATION 5.4 Hedging with two risk factors.

Suppose you manage a fund that invests primarily in oil refining stocks. As such, you are exposed to both movements in oil prices and in the stock market. Given the current uneasiness in the stock market, you find yourself in a dilemma. On one hand, you believe that there is a strong chance that the market will fall over the next couple of weeks due to a rise in the price of crude oil, but, on the other, that your particular portfolio of oil stocks will appreciate in value relative to the stock market due to the rising price of crude. Consequently, you want to hedge your market risk exposure, but not your oil risk exposure. Compute the stock market risk-minimizing hedge using the return data provided in ***Oil hedge.xls****. Use the S&P 500 futures contract traded on the Chicago Mercantile Exchange to represent the equity risk factor and the crude oil futures contract traded on the New York Mercantile Exchange to represent the oil risk factor. The contract denomination of the S&P 500 futures is 250 times the index level, and the denomination of the crude oil futures is 1,000 barrels.*

An important first step in an analysis of the hedge involving multiple risk factors is to understand the correlation among the returns series. The raw data in the file appears as follows:

Date	Oil Stock Portfolio Value	Mar. 2005 S&P 500 Futures Price	Mar. 2005 Crude Futures Price	Oil Stock Portfolio Return	Mar. 2005 S&P 500 Futures Return	Mar. 2005 Crude Futures Return
20040701	44,590,000	1128.50	37.00			
20040702	44,720,000	1127.30	36.64	0.00291	–0.00106	–0.00978
20040706	45,100,000	1116.70	37.63	0.00846	–0.00945	0.02666
20040707	45,370,000	1119.70	37.11	0.00597	0.00268	–0.01392
...	...	...	...	...	...	...
20050216	58,480,000	1210.50	48.33	0.02704	–0.00017	0.02239
20050217	58,130,000	1201.00	47.54	–0.00600	–0.00788	–0.01648
20050218	59,410,000	1202.30	48.35	0.02178	0.00108	0.01689
20050222	58,250,000	1184.70	51.15	–0.01972	–0.01475	0.05630

In Excel, go to the "Tools" menu, choose "Data analysis," and then "Correlation." This tool will allow you to generate the following matrix of correlations.

	Portfolio	S&P 500	Crude
Portfolio	1		
S&P 500	0.4641	1	
Crude	0.2738	–0.2060	1

The portfolio returns are strongly positively correlated with both the S&P 500 index, 0.4641, and the return of crude oil, 0.2738. At the same time, the S&P 500 return is inversely correlated with the return of crude oil, –0.2060. In other words, where the stocks, in general, fall as the price of crude oil rises, your particular portfolio of oil stocks tends to rise as crude oil rises.

Your objective is to hedge the stock market risk of your portfolio over the short-term. In order to estimate the appropriate hedge, you need to regress your portfolio returns on the returns of *all* known risk factors—in this case, the S&P 500 return and the crude oil return. The estimation results are as follows:

Regression Statistics	
Multiple *R*	0.5982
R-square	0.3579
Adjusted *R*-square	0.3498
Standard error	0.0080
Observations	162

	Coefficients	**Std. Error**	***t* Stat**
Intercept	0.0010	0.0006	1.6483
S&P 500	0.8252	0.0986	8.3697
Crude oil	0.1792	0.0302	5.9395

The regression results, indeed, confirm that the value of your portfolio increases with the stock market and crude oil. The estimated coefficient on the S&P 500 is 0.8252, which means that for a 1% change in the price of the S&P 500 futures contract, your portfolio increases in value by 0.8252%, holding the effects of crude oil constant.[10] The number of S&P 500 futures to sell is determined by using the estimated slope coefficient, the market value of the portfolio, the market price of the futures, and the futures contract denomination is the following way:

$$n_F = -0.8252\left(\frac{58{,}250{,}000}{1184.70 \times 250}\right) = -162.30$$

[10] This is precisely why it is important to include all possible risk factors. Otherwise, it is impossible to disentangle the effects of the different factors, except in the unusual case where the risk factors are independent of one another. In OLS regression, this problem is referred to as *omitting relevant explanatory variables*, and its consequences are discussed in Appendix B to this book.

Finally, in the interest of completeness, suppose we set the stock market risk-minimizing hedge using the estimated slope coefficient from a simply regression of portfolio return on stock index futures return. The regression results are:

Regression Statistics	
Multiple *R*	0.4641
R-square	0.2154
Adjusted *R*-square	0.2105
Standard error	0.0088
Observations	162

	Coefficients	**Std. Error**	***t* Stat**
Intercept	0.0014	0.0007	2.0722
S&P 500	0.7046	0.1063	6.6272

With the estimated slope coefficient being 0.7046, the number of S&P 500 futures contracts to sell is now

$$n_F = -0.7046\left(\frac{58{,}250{,}000}{1184.70 \times 250}\right) = -138.57$$

Selling this "risk-minimizing" number of futures is *wrong*. This number of contracts is *downward biased* because we failed to account for the fact that the crude oil return and the S&P 500 return are negatively correlated. Without the crude oil return in the regression, the S&P 500 return is proxying for two factors—the stock market and crude oil. Since the crude oil and the stock market are negatively correlated, this means that the slope coefficient in the simple regression on the S&P 500 futures return is downward biased. If the correlation had been positive, the slope coefficient in the simple regression would have been upward biased.

Finally, with multiple risk factors, measuring the effectiveness of the hedge becomes slightly more complicated. The *R*-squared in the first regression, 0.3579, says that 35.79% of the variance of the return of the oil stock portfolio can be explained by the returns of the S&P 500 futures and the crude oil futures. But, what percentage of this risk remains after the S&P 500 futures hedge is in place?

To answer this question, we compute the standard deviation of the unhedged portfolio return, $\sigma(\tilde{R}_P)$, as well as the standard deviation of the hedged portfolio return, $\sigma(\tilde{R}_V - 0.8252 \times \tilde{R}_{S\&P500})$. They are 0.00993 and 0.00540, respectively. Thus the proportion of the unhedged portfolio return variance that remains after the hedge is put in place is

$$\frac{0.00883^2}{0.00993^2} = 79.09\%$$

ESTIMATION ISSUES

There are some subtle regression estimation issues that worth noting at this juncture. An important one is the proper selection of the frequency of the price observations used in generating price changes or returns. The regression can be performed on daily, weekly, or monthly price changes or returns—given that price changes/returns are i.i.d., it should not matter. From a purely statistical standpoint, however, the higher is the frequency, the better. The greater is the number of observations for a given historical time period, the greater the amount of information that gets impounded in the estimate. From a practical perspective, however, there is a tradeoff. While greater frequency means more information, it also means more measurement error.

Measurement errors arise from a variety of sources. We will discuss three—bid/ask price bounce, nonsimultaneous price observations, and infrequent trading. Before addressing the effects of these potential sources of error, it is useful to think conceptually about the use of regression analysis for setting hedge ratios. Implicitly or explicitly, we made a number of assumptions. A critically important one was that the relation between asset returns (price changes) and futures returns (price changes) was *stationary* through time. Among other things, this allowed us to project expected *future* variances and covariances that go into setting the hedge ratio from *past* price data. Other assumptions are also critical. In using historical price data, we implicitly assume that we are measuring "true" prices[11] and that the asset and futures prices at a given time t are observed at exactly the same instant in time. Both of these assumptions are normally violated. Indeed the magnitude of the errors induced by these considerations may be quite large.

Bid/Ask Price Bounce

The problem here emanates from the fact that the daily prices recorded in historical data bases are usually last trade prices. A *last trade price* is the price recorded at the time of the last transaction of the day. In general, the last trade price will *not* be the security's true price. One reason for this is that, in all likelihood, the trade took place on one side of the market. If the last trade was seller-motivated, the trade was probably consummated at the prevailing bid price, and, if the trade was buyer-motivated, it probably took place at the ask. (See Figure 5.4.) A better proxy for the end-of-day true price is the midpoint of the prevailing bid/ask price quotes at the end of the day, however, histories of daily price quotes are not generally available.

Now consider how prices are used in the regression of asset return on futures return. For the asset, return is measure from close to close. The computed asset return, therefore, has two measurement errors, one for each price. The same is true for the futures. Thus, the number of measurement errors included in a regression using T days of returns is $4 \times (T + 1)$. This *errors-in-the-*

[11] Here we are considering the "true" price of the security to be its price in a frictionless market.

FIGURE 5.4 Relation between last trade price of the day and true price.

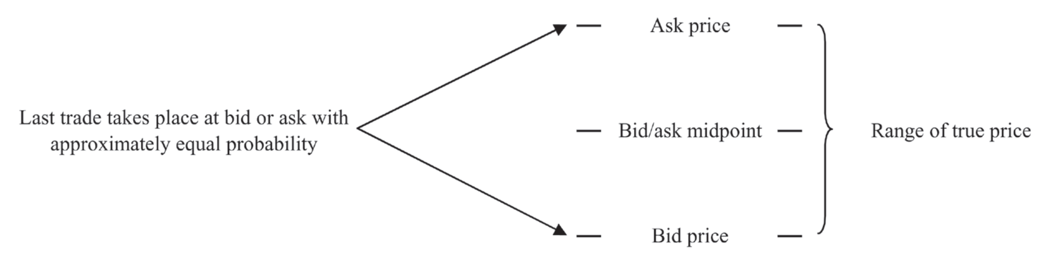

variables problem tends to bias the estimated slope coefficient downward.[12] Note that, if the spreads between the bid and ask prices in the asset and futures markets is zero, this problem disappears. Conversely, in markets that are not particularly liquid, bid/ask spreads will be high and the effects of this problem will be large. Note also that we can control for the bias induced by bid/ask price bounce somewhat by lengthening the period over which returns are measured. The reason is that we are reducing the sheer number of errors. In regressions based on true prices, daily data are expected to produce the same slope coefficient as weekly or monthly data. A regression involving weekly returns based on last trade prices, however, will have proportionately fewer errors than a daily return regression (with five trading days per week, about five times fewer errors), and, hence, will produce a slope coefficient that is less downward biased.

Nonsimultaneous Price Observations

Another source of measurement error is that the price observations that are used in generating returns for the asset and the futures may not be simultaneous. Consider regressing the daily returns of a stock portfolio on the returns of the S&P 500 futures contract. If we use closing prices for each market, we have another errors-in the-variables problem because the stock market closes at 4:00 PM (EST), while the S&P 500 index futures market closes at 4:15 PM. This timing mismatch causes that the slope coefficient in the regression to be downward biased. Part of the observed futures return will not be reflected in the stock portfolio return until the following day.[13]

Infrequent Trading

Yet another problem arises when the asset that we are trying to hedge is an amalgam of other asset prices that have varying degrees of trading frequency. Consider the closing index level of the S&P 500 portfolio each day, for example. The "true" S&P 500 index level at the close should be based on a weighted-average of the "true" prices of all 500 index stocks, where each and

[12] For a discussion of the errors-in-the-variables problem, see Appendix B of this book.

[13] In principle, this problem could be handled by including a lagged futures return in the regression model. The procedures for doing so are somewhat inexact, however.

every stock traded at exactly the close. But each stock did not trade at the close. While in a typical trading day stocks like General Electric and IBM trade almost continuously right up until the close, others trade fairly much less frequently. Indeed some stocks may not to have been traded during the last few hours of the day. For these stocks, the last trade price is a poor indicator of the true end-of-day price. Indeed, because the index is computed on the basis of last trade prices of the constituent stocks, the "observed" index will always lag its true level. Among other things, this means the observed index returns will be positively serially correlated, thereby violating the regression assumption that returns are i.i.d.

ILLUSTRATION 5.5 Testing for robustness.

*The file **High yield.xls** contains the daily returns of the Merrill Lynch high-yield B bond index, the CBT's T-bond futures, and the CME's S&P 500 futures. Bonds rated below BBB are sometimes called "junk" bonds. Junk bond prices are usually sensitive to both long-term interest rate and stock market movements. Regress the daily return of the bond portfolio on the returns of the T-bond futures and S&P 500 futures as if you were attempting to hedge the bond portfolio risk factors. Now create Wednesday to Wednesday returns by summing the daily returns over the week, and run the same regression on weekly returns. Compare and comment on the regression results.*

The table below summarizes the regression results. The daily regression results indicate that the bond portfolio value is relatively insensitive to movements in the T-bond futures price (long-term interest rates) and in the S&P 500 index. For a 1% change in the T-bond futures price, the bond portfolio value changes by 0.0693%, and a 1% change in the S&P 500 futures price causes the bond portfolio value to change by 0.0138%. The hedging effectiveness appears to be very low, at 11.93%.

The weekly results are quite different. The coefficient estimate on the T-bond futures return is 0.1625 and on the S&P 500 futures return is 0.0440, both more than twice as high as in the daily regression. In addition, the hedging effectiveness is more than twice what it appeared in the daily regression. The biweekly results improve matters even further.

		α_0	α_1	α_2	R^2
Daily	Coefficient	0.0003	0.0693	0.0444	0.1193
	Std. error	0.0001	0.0138	0.0118	
Weekly	Coefficient	0.0014	0.1625	0.1401	0.2668
	Std. error	0.0006	0.0440	0.0416	
Biweekly	Coefficient	0.0026	0.1758	0.1725	0.3546
	Std. error	0.0012	0.0590	0.0563	

These results reveal the danger in blindly applying regression analysis. The nature of the data used in the regression needs to be carefully considered before the regression is performed. In this particular instance, the most likely culprit is the bond index. Corporate bonds with a high degree of credit risk trade infrequently. Indeed, it is not uncommon for some corporate bonds to be traded shortly after issuance and then never again. Consequently, it is highly likely that the bond index suffers from infrequent trading effects.

One way to test this possibility is to examine the autocorrelation function of the bond portfolio returns. This can be done using the OPTVAL Library function

OV_STAT_AUTOCORREL(k, x, *out*)

where k is the maximum number of lags, x is the time series vector, and *out* is an indictor variable controlling the output vector (0 = horizontal, and 1 = vertical).[14] With 294 daily returns in the time series, the standard error is $1/\sqrt{294} = 0.58$.

Lag	1	2	3	4	5
Bond portfolio	0.5578	0.3684	0.3232	0.2175	0.1772
T-bond futures	0.0206	–0.0941	0.0492	0.0202	0.0201
S&P 500 futures	0.0287	–0.0542	–0.0147	–0.0556	0.0349

As the results show, we can reject the hypothesis that the first-order autocorrelation of the daily bond portfolio returns is zero. At the same time, we cannot reject the hypotheses that the T-bond futures returns and the S&P 500 futures returns are uncorrelated.

Like the bid/ask price effects, the effects of infrequent trading begin to disappear as the distance between adjacent price observations used in the computation of returns becomes larger. The autocorrelation functions for the weekly and biweekly bond returns are as follows:

Lag	1	2	3	4	5
Weekly	0.3148	–0.0804	0.0177	–0.0411	–0.1000
Biweekly	0.1219	–0.1179	0.0759	0.0245	–0.0528

where the standard error for the weekly correlations is 0.130, and the standard error for the biweekly correlations is 0.183. What the results indicate is that the first-order autocorrelation remains significant in the weekly returns but disappears for biweekly returns. Based on this analysis, it is safer to use the regression results from the biweekly regression in setting the risk-minimizing hedge.

SUMMARY

This chapter explores the role of forward/futures contracts in managing expected return and risk. We begin by showing how futures contracts fit within the CAPM and develop the concept of hedging. We then focus on using futures to manage different types of risks. We start with price risk and consider the case in which an airline wants to hedge the cost of jet fuel. We then focus on revenue risk and show how a farmer can hedge the sales proceeds of his crop in an environment with both price and quantity risks. Next we consider gross margin (i.e., uncertain revenue less uncertain costs) risk. Oil refiners, for example, are concerned about the difference between the revenue they realize through the sale of heating oil and unleaded gasoline and the cost of the crude oil they must acquire

[14] The use of the autocorrelation function is described in greater detail in Appendix A of this book.

to produce these products. Finally, we consider risk management when the asset or portfolio has multiple risk factors. Someone managing a junk bond portfolio, for example, faces both interest rate and stock market risk exposures. For the most part, the illustrations discussed in this chapter are confined to risk-minimizing hedges. The principles can easily be extended to include the tradeoff between expected return and risk. For setting risk-minimizing hedges, OLS regression winds up being an indispensable tool. As important as the regression technique, however, is a thorough understanding of the data being used in the regression estimation.

REFERENCES AND SUGGESTED READINGS

Anderson, Ronald W., and Jean-Pierre Danthine. 1980. Hedging and joint production: Theory and illustrations. *Journal of Finance* 35: 487–498.

Anderson, Ronald W., and Jean-Pierre Danthine. 1981. Cross hedging. *Journal of Political Economy* 89 (December): 1182–1196.

Black, Fischer. 1976. The pricing of commodity contracts. *Journal of Financial Economics* 3 (March): 167–179.

Dale, Charles. 1981. The hedging effectiveness of currency futures markets. *Journal of Futures Markets* 1: 77–88.

Eaker, Mark, and Dwight Grant. 1990. Currency hedging strategies for international diversified equity portfolios. *Journal of Portfolio Management* 17 (Fall): 30–32.

Ederington, Louis H. 1979. The hedging performance of the new futures markets. *Journal of Finance* 34 (March): 157–170.

Figlewski, Stephen. 1984. Hedging performance and basis risk in stock index futures. *Journal of Finance* 39 (July): 657–669.

Figlewski, Stephen, Yoram Landskroner, and William L. Silber. 1991. Tailing the hedge: Why and how. *Journal of Futures Markets* 11: 201–212.

Frankle, C. 1980. The hedging performance of the new futures markets: Comment. *Journal of Finance* 35: 1273–1279.

Grant, Dwight. 1985. Theory of the firm with joint price and output uncertainty and a forward market. *American Journal of Agricultural Economics* 67: 630–635.

Howard, Charles T., and Louis J. D'Antonio. 1994. The cost of hedging and the optimal hedge ratio. *Journal of Futures Markets* 14 (April): 237–258.

Park, T., and L. Switzer. 1995. Bivariate GARCH estimation of the optimal hedge ratios for stock index futures: A note. *Journal of Futures Markets* 15: 61–67.

Working, Holbrook. 1953. Futures trading and hedging. *American Economic Review* 48: 314–343.