

Philip Dybvig is the Boatmen's Bancshares Professor of Banking and Finance at Washington University in St. Louis. This article was prepared while he was a consultant at the Research Department of the Federal Reserve Bank of St. Louis. William J. Marshall is the Chief Operating Officer of NISA Investment Advisors, L.L.C. The authors are grateful for useful suggestions from Bill Gavin, Myung-Yull Pang, and Murray Weidenbaum.



The New Risk Management: The Good, the Bad, and the Ugly

Philip H. Dybvig and
William J. Marshall

At one time, risk management meant buying corporate insurance, implementing procedures to avoid lawsuits and accidents, and installing safety equipment. The new risk management uses financial markets to hedge different sources of risk within the firm. Trading in financial markets can hedge companies against the risk of changes in interest rates, input prices, or currency fluctuations. While hedging per se is not new, the scale and diversity of hedging are far greater than they used to be. When executed properly, the new risk management can be good and even essential for competition. Unfortunately, the new risk management can also be bad, wasting resources without reducing risk and perhaps even increasing it. The new risk management can be ugly, generating large losses such as those in widely publicized cases at Barings, Metallgesellschaft, Procter and Gamble, and other firms. In these and many other firms, employees relatively far from the top of the hierarchy of control had the authority to take financial positions large enough to generate losses that could bankrupt the firm. Thus, policies for risk management should be put in place at the highest level of a firm, and they should provide for monitoring and control. The purpose of this article is to provide an introduction to the new risk management

and some policy choices firms should be considering.

We start with a discussion of the option-pricing tools that make the new risk management possible, and we follow with a stylized example of how the new risk management ought to work. Then we consider implementation issues, including some general policy questions as well as some accounting issues.

TOOLS FOR THE NEW RISK MANAGEMENT

Starting with the famous work of Black and Scholes (see shaded insert, next page), option-pricing theory has been very successful in pricing various financial claims. The Black-Scholes model was designed to price standard call and put options, and it has been extended to price all sorts of financial claims. The Black-Scholes model and its extensions form the theoretical foundation for the new risk management.

There were option-pricing models prior to the work of Black and Scholes, including some models with formulas similar to Black-Scholes. What makes the Black-Scholes model different is that it provides a *hedging strategy* that is an investment policy with an investment equal to the model's option price and a terminal value equal to the terminal value of the option. Knowing the trading strategy means that the model is not only someone's best guess; it is also possible to profit if the model is wrong. If the model price is lower than the price in the economy, we can sell the option, pocket the excess over the model price, and invest in the hedging strategy to cover the terminal value of the option we have sold. If the model price is higher than the price in the economy, we follow the hedging strategy in reverse, taking a short position instead of a long position and lending instead of borrowing. In the model, the

THE BLACK-SCHOLES OPTION-PRICING MODEL

The precursor of all modern option-pricing models was developed by Fischer Black and Myron Scholes.¹ The main result is an option-pricing formula based on simple and reasonable assumptions in a continuous-time model. The remarkable thing about the result is that it relies on the absence of arbitrage, and part of the proof is a formula that specifies a trading strategy in the underlying stock and the riskless bond that will replicate the payoff of the option at the end.¹¹ If the option is priced differently in the economy, buying or selling the option and following either the trading strategy or the reverse of the trading strategy will make money! Using the same sort of analysis, one can derive a trading strategy that will hedge the financial risk in a firm's cash flows.

Now we present the Black-Scholes formula for the price of a call option. Recall that a call option gives the owner the right (at the owner's option) but not the obligation to buy one share of the underlying stock at the strike (or exercise) price X specified in the option contract on or before the maturity date of the option. If the stock price is S and the price of the bond promising to pay the amount of the strike price at the maturity date of the option is B , the Black-Scholes price, C , of the call option is

$$C = S N(x_1) - B N(x_2),$$

where

$$x_1 = \log(S/B)/s + s/2,$$

$$x_2 = \log(S/B)/s - s/2, \text{ and}$$

s is the standard deviation (or square root of the variance) of the stock price at maturity, given the stock price today, and the function $N()$ is the cumulative normal

hedge replicates the option value perfectly; in practice, the hedge is not perfect, but it works remarkably well. This is why the Black-Scholes model and its progeny are widely used in business.

The introduction of these option-pricing models and the parallel development and maturation of liquid financial markets have made it easier and easier to hedge financial risks using options, futures, futures options, swaps, caps, collars, floors, and a variety of other financial instruments.

OPTION PRICING AND RISK MANAGEMENT

Hedging an option is an example of risk management. Its purpose is to remove the risk and capture the pure economic profit of the transaction. Fundamentally,

this strategy is the same as insurance. For the insured, the insurance policy makes money in bad times (when the insurable event occurs) and loses money in good times (when no insurable event occurs but the premium is paid), which reduces risk by softening the impact of bad outcomes. The same is true of a hedging strategy; losing money on the hedge in good times and making money in bad times offsets the original cash flows, making the total cash flow less volatile. In either strategy, payment for the insurance can be "up-front" or "pay-as-you-go": For hedging, as for insurance, the arrangement of cash flows¹ accommodates the preference of the insured. There are important differences in taxation and regulation between hedging using insurance and hedging using financial markets, but those are beyond the scope of this paper.

¹ Cash-flow is the accounting notion of actual cash coming in or out from operations. Unlike profits, cash flow does not include depreciation or amortization, but it does include (as an offset) investment in capital. In our examples later, we will treat the two the same, although this is not appropriate except in the case of very simple businesses that rent any required capital.

distribution function. If there is a constant, continuously-compounded interest rate, r , and T is the time-to-maturity of the option, then B is the discounted exercise price

$$B = X \exp(-rT).$$

And, if the stock has a variance, v , per unit time, we have that

$$s^2 = vT$$

is the variance of the final stock price.

In the expression for C , the first term is the stock holding in the hedge strategy, and the second term is the bond holding (which is negative, which is a short sale or borrowing). The main assumptions of the model are absence of arbitrage, a constant riskless rate, continuous stock prices, and a constant variance of returns per unit of time for the underlying stock. The intuition is that we can replicate the risk of holding the option by holding just the right portfolio of riskless bonds and the underlying stock. For example, if at a point in time the option moves fifty cents for each one-dollar movement in the underlying stock price, then the replicating strategy would hold one share of stock for each two options we are replicating. To hedge the value of the option, we would short (borrow) a share of stock for each of two options. In that case, the stock's value change would neutralize the effect on our wealth of the option's price change. The hedge's holdings in the stock and bond will change over time and in response to stock price changes, since the sensitivity of the option value is different when the option is in the money than when it is out.

[†] Black, F. and M. Scholes, "The Pricing of Options and Corporate Liabilities" *Journal of Political Economy* 81, 1973, 637-54.

^{††} For more discussion of why this makes sense, see Rubinstein, M. and H. Leland, "Replicating Options with Positions in Stock and Cash," *Financial Analysts Journal*, (Jan-Feb 1995), pp. 113-21.

Using dynamic trading strategies to hedge financial options may seem significantly different from hedging price risk in a firm. However, the concept is exactly the same. A hedger is taking the other side of the risky investment in futures or whatever would be used to replicate the cash flows that are being hedged. Normally, these cash flows cannot be hedged precisely, but the hedge can still reduce risk significantly. For example, one policy is to hedge the expected cash flow conditional on the price of inputs that can be hedged in futures markets while leaving the remainder unhedged, which means that the remaining risk is borne by the stock and bond holders of the firm.

Before turning to the general policy issues in risk management, we will consider a typical example.

RISK MANAGEMENT IN MANUFACTURING

Our example considers the hedging problem of a manufacturer that uses significant amounts of copper as an input. (With little change in the discussion, this input could be zinc, silver, oil, or wheat. With a slightly greater change, the "production" could be servicing of core deposits in a bank, and the analysis would provide the optimal hedging of interest rates.) We will examine the optimal hedging of copper price movements in the cash flows before turning to a general discussion of policy and oversight.

In the example, expected output is 1,000 units, which will sell for \$100 per unit. The price has been committed to in advance because of long-term contracts,

Table 1

A Manufacturer's Copper Price Hedge

Each section of this table shows the cash flows one year from now for the simple example of a manufacturer that is facing copper price risk. In the example, copper prices are higher when demand for output is higher. Each section of the table illustrates a different hedging strategy and profit (= cash flow) in three copper price scenarios. The example abstracts from taxes and sources of risk that are not related to the price of copper. In each case, the expected profit is 1,500. The point of hedging is reducing uncertainty, not increasing average cash flow (except indirectly, because it allows you more freedom in choosing projects).

Table 1A

Unhedged Cash Flows

Probability	Copper Price	Units Sold	Output Price	Total Sales	Copper Expense	Other Expenses	Profit (Loss)
1/4	25	1,200	100	120,000	30,000	82,000	8,000
1/2	20	1,000	100	100,000	20,000	78,000	2,000
1/4	15	800	100	80,000	12,000	74,000	(6,000)

Table 1B

Naive Hedge of the Expected Quantity Required

This hedge might be put in place as part of the procurement process, since it looks only at expenses. This is at best an incomplete hedge of copper costs, since the true quantity changes with copper prices. In our example, this naive hedge actually increases risk, since increased sales mean profits are high when copper prices are high.

Probability	Unhedged	Hedge	Net
1/4	8,000	5,000	13,000
1/2	2,000	0	2,000
1/4	(6,000)	(5,000)	(11,000)

Table 1C

Fully Hedged Cash Flows

A complete hedge of all the cash flows requires something more than a simple purchase of futures, since the sensitivity to copper prices of the unhedged profit or loss is higher when copper prices are low than when copper prices are high.

Probability	Unhedged	Hedge	Net
1/4	8,000	(6,500)	1,500
1/2	2,000	(500)	1,500
1/4	(6,000)	7,500	1,500

but the quantity may vary around this expectation because the contracts give customers the option to choose how much to buy within a range. Each unit will use an amount of copper that would cost \$20 purchased forward (in a firm commitment to buy one year from now). If purchased in

the spot market, the copper in the unit might cost \$25 (with probability 1/4), \$20 (with probability 1/2), or \$15 (with probability 1/4).

One obvious (and common) approach to hedging in this context would be to forecast demand for copper and then hedge that amount, either by entering a fixed-price contract with the supplier or by buying that amount of copper futures—at a shorter maturity, if necessary, because one-year futures are not traded or have a very large spread. This might be a natural outcome if hedging were performed by buyers who were responsible for copper procurement and whose evaluations were based on the cost of a forecast quantity of copper. However, choosing a useful hedge of the entire cash flow is more subtle than that.

Table 1 contains an elaboration of the example. When the economy is doing well, copper prices are high (since this firm and other manufacturers are demanding more copper) and so is demand for the firm's output. Table 1A shows the cash flows in the absence of any special risk management to hedge copper price risk. Table 1B shows the result of hedging by buying forward the expected quantity. Ironically, this naive approach to hedging increases risk exposure, since the firm is already more than hedged by increased sales when the industry is doing well and copper prices rise. The full hedge,

the result of which is shown in Table 1C, cannot be implemented by simply buying or selling copper forward one year. However, the full hedge can be implemented either by buying options or by dynamic trading in forward or futures contracts. Since this type of strategy is typical of hedging problems, it is worthwhile deriving the dynamic hedge and discussing its operation.


THE DYNAMIC HEDGE

To study the dynamic hedge, we need to understand the trading opportunities and information between now and realization of the cash flows a year from now. The sensitivity of the firm's value to copper prices varies in response to the interim information, and this changing sensitivity should be reflected in our trades.

In the current example, we assume that the firm is using copper futures contracts to hedge changes in copper prices. Futures serve the same economic purpose as forward purchases, but they are somewhat different logistically, since for futures the money changes hands immediately when the prospective value of copper rises and falls. If we buy one futures contract, then at the end of each day we are given (more literally, our margin account is credited with) the change in futures price over the day. If we sell (or *short*) one futures contract, then we must pay the change. If the futures price increases from \$50 to \$55, then the owner of two futures contracts will collect \$10, and someone who has sold two futures contracts will have to pay \$10. If the futures price instead decreases from \$50 to \$45, the person short two contracts collects \$10, and the person long two contracts has to pay \$10. In general, the futures price need not be exactly equal to the price we would pay for forward purchase, but for most purposes we can think of the two as being the same.²

In the actual economy, information arrives minute-by-minute, and a firm can trade on copper prices almost continuously in time. For our simple example, informa-

Figure 1



tion arrival and trading occur now, six months in, and again in a year. (This is not an essential simplification; while the analysis for a practical model requires more computations, it is conceptually no more difficult.) At the beginning of the year, the futures price of copper delivered a year from now is \$20. Six months from now, the futures price will be either \$22.50, with probability 1/2, or \$17.50, also with probability 1/2. The overall price dynamic is given in Figure 1. The price at a node in the tree is the price paid in a firm commitment to buy copper one year from now. From a given node, an up or down move is equally likely, with probability 1/2, so any given price path has probability $1/4 = 1/2 \times 1/2$. Consistent with Table 1, the ending node of \$20 is twice as likely as the other ending nodes because it can be reached by either an up move followed by a down move (probability 1/4) or a down move followed by an up move (probability 1/4). A final price of \$25 comes only from two up moves (probability 1/4), and a final price of \$15 comes only from two down moves (probability 1/4).

To derive the full dynamic hedge, the firm requires one more piece of information, which is the rate at which futures gains or losses will be reinvested, which we will take to be 5 percent simple interest over six months. (Actually, the rate we choose will not affect the hedged cash flows in Table 1C, since increasing this rate will result in a completely offsetting decrease in the number of contracts we hold over the first six months.) Holding one futures contract at one node implies a gain of \$2.50 (given an up move) or a loss of \$2.50 (given a down move), which

² In fact, if interest rates are non-random (so re-investment rates are known in advance), absence of arbitrage implies that the forward price must equal the futures price, although one futures contract has more impact, since the change in value is received up front, while in a forward contract the change in value occurs at maturity.

A SMALL GLOSSARY OF RISK-MANAGEMENT TERMS

Binomial model. The binomial option-pricing model, developed by Cox, Ross, and Rubinstein [1979], assumes that the stock return over a short time interval has one of two values. The binomial model is a popular alternative to the Black-Scholes model because it is flexible and easy to implement on a computer.

Black-Scholes model. This is the original modern option-pricing model (see shaded insert on pp. 10-11).

Call option. A call option is a contract that gives the owner the right to purchase a share of the underlying asset in exchange for the contractually specified strike price (or exercise price). An *American* call option can be exercised at any time before maturity, while a *European* call option can be exercised only on the maturity date.

Cap. An interest-rate cap is a promise to pay the excess of an interest rate above some level in each of a number of periods. Caps are useful for containing the risk of rising borrowing costs.

Collar. A collar combines the cash flows of buying a cap and selling a floor. It is useful for containing the risk of rising interest rates (like a cap); including the floor gives up some profit potential when rates fall to help to pay for the cap.

Floor. An interest-rate floor is a promise to pay the shortfall of an interest rate below some level in each of a number of periods. Floors are useful for locking in a minimum return.

Forward contract. A forward contract gives the owner the right and the obligation to buy a specified amount of a commodity at a specified price at some specified date in the future.

Futures contract. A futures contract is similar to a forward contract except that there is daily settlement, i.e., each day the parties to the contract exchange money representing the market-determined change in value of the contract. Daily settlement minimizes the need for credit checks and large margin accounts (which are held as collateral), since only one day's price variation is at risk.

Hedge. Hedging a position (or entering a hedge) is undertaking another activity with offsetting risk. Some common hedging instruments include insurance, futures contracts, and options.

Long position. To take a long position (or to "be long") is to purchase an asset or futures.

Put option. A put option is a contract that gives the owner the right to sell a share of the underlying asset in exchange for the contractually specified strike price (or exercise price). An *American* put option can be exercised at any time before maturity, while a *European* put option can be exercised only on the maturity date.

Short position. To take a short position (or to "sell short") is to assume the opposite of a long position. In the case of futures, they are simply sold in the market. Shares and other securities are borrowed (for a nominal fee) then sold in the market, with the promise of buying some shares later to return the borrowed shares. In the meantime, the short must pay any dividends or coupons that are due the person from whom the shares were borrowed. The cash flows for a short position are the negative of the cash flows for a long position.

Value at risk. Value at risk (VAR) is a measure and methodology for assessing risk exposure by looking at total exposure to various market-level risks. This is a useful tool, but it does not account for residual risk that is specific to the project and not related to the market-level risks.

Table 2

Cash Flows from the Dynamic Hedge

Futures Price Path	Cash Now	# Contract Now	Cash in 6 months	# Contracts in 6 months	Cash in 1 yr	Pre-hedge Cash Flow	Hedged Cash Flow
20-22.50-25	0	(1,333)	(3,333)	(1,200)	(6,500)	8,000	1,500
20-22.50-20	0	(1,333)	(3,333)	(1,200)	(500)	2,000	1,500
20-17.50-20	0	(1,333)	3,333	(1,600)	(500)	2,000	1,500
20-17.50-15	0	(1,333)	3,333	(1,600)	7,500	(6,000)	1,500

is reinvested until the end at the interest rate. From this we can use simple algebra to derive the solution. In the example, the full hedge is implemented by the following strategy: At the start, the firm sells $1400/1.05 \sim 1,333$ futures at the futures price of \$20. If futures go down to \$17.50, the firm increases the short position to 1,600 contracts, while if futures go up to \$22.50, the firm reduces the short position to 1,200 contracts.

The terminal cash flow generated by the hedge (including reinvestment) is analyzed in Table 2. For example, the second row shows the effects of the hedge when prices go up and then down (from \$20 to \$22.50 and back to \$20). The hedge starts with no initial cash. It shorts 1,333 contracts, and when in six months the futures price goes up by \$2.50, $\$2.50 \times 1,333 = \$3,333$ is borrowed, and the short futures position is reduced to 1,200 contracts. When the futures price falls by \$2.50, $\$2.50 \times 1,200 = \$3,000$ in profits are collected, and after payment of $\$3,333 \times 1.05 = \$3,500$ on the loan, net cash from the hedge is $\$3,000 - \$3,500$ for a loss of \$500. Added to the unhedged cash flow in that state of \$2,000 (from Table 1A), the hedged cash flow is \$1,500. The calculations in the other states work the same way.

We can see now that the dynamic hedge was chosen so that the re-invested proceeds of the hedge plus the original cash flows are made to be the same in every contingency. The necessary hedge can be computed by working backward from the end. The first two rows differ only in the price performance over the last period. Since the difference in pre-hedge cash flow for these two scenarios

is $\$8,000 - \$2,000 = \$6,000$ and the difference in futures prices for the two scenarios is $\$25 - \20 , we require $\$6,000 / \$5 = 1,200$ contracts to replicate the cash flows or short 1,200 contracts (the offsetting position) to hedge the cash flows. Given the calculated hedge at the last date, the calculation at the next earlier date proceeds in the same way, and so forth back to the start. The entire strategy can be computed by looking at the linear equations implicit in Table 2, or by standard techniques described in option pricing textbooks.

While the model underlying the hedge for the simple example probably seems too simple, it is in fact similar (except for the number of intermediate trading dates) to the binomial models used successfully in practice. Adding the additional subperiods is straightforward, given modern computing resources.

SOME FUNDAMENTAL QUESTIONS

In the example in the previous sections, we assumed that hedging is desirable. However, this assumption is far from obvious, and it is useful to examine potential motives for hedging.

Why Should We Hedge?

The reason for hedging should link back to the overall objective of the firm, which is to create or enhance economic value. There is a general issue of whether the firm should maximize narrowly the value to shareholders, the total value to all financial claimants, or some more general

social value to a variety of stakeholders. This distinction will not be so important to us; most importantly, we will assume that taxes (governments' claims) are not part of what we are optimizing, and for concreteness we will speak of maximizing value to shareholders in the firm.

The first and most obvious effect of hedging is that it reduces the volatility of the value received by shareholders. Unfortunately, this does not have any value for most shareholders in a large publicly-traded firm, who hold the shares in a well-diversified portfolio and for whom the additional risk is unimportant. Indeed, a conflict of interest may exist between the majority of shareholders and large shareholders (for example, members of the founding family who hold 30 percent of the shares and whose holdings are undiversified): expending resources to reduce risk may benefit the large shareholders at the expense of the rest of the shareholders. Management may have a similar conflict, since risk threatens their jobs and they may have a significant proportion of their wealth tied up in the firm's shares. Since most shareholders in a publicly traded firm would not care about the additional risk attributable to copper price exposure, this is not a good reason for hedging. (On the other side of the equation, the cost of hedging may be very small; we will consider this consideration further in a later section on cost issues.)

A more subtle argument for managing copper price risk is that failure to do so may cause ancillary damage within the firm. As an extreme case, adverse copper price movements may push the firm into bankruptcy, which has a number of deadweight costs to the firm, such as payments to lawyers and accountants and the loss of profitable future projects. More normally, unhedged risk exposure may tend to increase taxes, on average: While the government receives additional tax payments when the copper price move is favorable, an unfavorable move will not create a compensating tax reduction, given that tax offsets may only be deferred (and may even be lost). A related tax reason for managing copper prices is

that the reduction of risk makes it possible to maintain more leverage to reduce corporate taxes and avoid "double taxation."

"Double taxation" is the payment of both corporate and personal taxes on cash flows going to equity, compared with payment of only personal taxes on cash flows going to debt, since interest expense is an offset to income in the computation of corporate income tax. While there are no personal taxes for institutional investors—and therefore no *double* taxation—the parallel argument—single taxation versus no taxation—is valid and even more powerful for institutions. For individuals there is at least a possibility that the corporate tax on equity will be offset by lower taxes at the individual level through deferred realization of gains or by a lower capital gains rate. For tax-exempt or tax-deferred investors, the extra tax is unmitigated.

A third argument for managing copper price risk is that many firms have a policy of smoothing earnings, and hedging can reduce volatility in earnings. Although this is common practice, it is hard to endorse, since it seems to be an expenditure of the owner's resources to minimize the amount of information getting out to the owners. (In principle, smoothing earnings might be used to eliminate temporary variations and provide a clearer picture of long-term value, but it seems more typical that smoothing is intended to avoid bad-looking quarters without necessarily distinguishing short- and long-term shocks.) This use of hedging may make management more comfortable and minimize criticism, but this is not obviously in the interest of shareholders. In some cases, hedging could be justified by the argument that it avoids restrictive debt covenants, but such covenants are far from binding in all but a small proportion of firms that smooth earnings. More common is the opposite extreme case, in which the internal objective of the firm is to ensure that earnings do not fall. Hedging for this purpose may make management comfortable—indeed too comfortable—but it discourages profitable innovation. A related strategy for keeping volatility of earnings small is to

maintain a low level of financial leverage, which implies a large voluntary tax contribution that is not in the interest of shareholders.

A fourth argument for managing copper price risk is to make it easier to give managers incentives to produce profits: By hedging risk, we can make (for example) a division manager's compensation depend closely on value-added that the manager can influence rather than what the manager can't influence (the actual realization of copper prices). This argument for managing copper price risk implies that it may be optimal to manage copper price risk at the division level even if copper prices do not represent a significant contribution to the firm's cash flow. Of course, this strategy begs the question of why it can't be done more cheaply (for compensation purposes only) using a paper portfolio.

What Risks Should We Hedge?

The question of what risks to hedge must be subordinate to the question of why we should hedge. If there is not a compelling reason to hedge a particular source of risk, then we probably should not be hedging it. One important issue is the sense in which we would hedge a certain type of risk. For example, suppose we are hedging a bank's exposure to interest-rate risk. Should we hedge the direct interest mismatch of existing assets and liabilities, or should we hedge the full economic value, which would include the value of future business? For example, a bank may find that, as interest rates rise, core deposits tend to be lost. Current accounting methods make it hard to hedge this sort of risk without penalty (and the risk-based capital requirements from the Basle Accord penalize almost all hedging because one has to increase capital once for the underlying cash flow and once again for the hedge). There is a related question of whether to hedge cash flows or value. In principle, the two are the same (if we were to hedge cash flows far enough out), but, in practice, hedging cash flows out a year is much different from

hedging the firm's entire value. If the purpose of hedging is to eliminate sources of noise that are beyond the manager's control, it may even be appropriate to hedge particular accounting numbers used in computing compensation rather than hedging cash flow or economic value.

With What Instruments Should We Hedge?

For most commonly hedged risks (such as exposure to interest rates, foreign exchange rates, or commodity prices), many instruments can be used for hedging. For example, to hedge U.S. interest rates we can use bonds, repurchase agreements, Treasury bond futures, swaps, caps, or collars. The choice among this set would be determined by pricing and transaction costs, match to hedging needs, and accounting implications.

Support Your Investment Banker

A common approach of managers planning to hedge is to turn the whole problem over to an investment banker who, after all, has the expertise and the traders who can put the hedge in place and is happy to provide "free" advice on what to do. As in all markets, the "free" advice is priced out in what you pay for the hedge, and then some. To avoid paying too much, it is best to understand how the hedge works and how much it should cost. Ideally, such expertise should be located in-shop; otherwise, it is worth the expense of hiring an expert to monitor the prices being paid to the investment banker. In general, competition among investment bankers may be useful in reducing the cost, but competition will not necessarily produce any incentive to report when hedging is unnecessary.

ACCOUNTANT: FRIEND OR FOE?

Suppose we put in place the optimal hedge computed above, using the model for demand and option-pricing theory to

determine the correct holding in futures to offset the risk in the cash flows. What will this do to our accounting statements?

In general, accounting looks at the present and the past: Accountants favor methods whose results are easily replicable, especially since standard mechanical rules, even if inaccurate, are easy to defend if the firm has followed Generally Accepted Accounting Principles (GAAP). Hedge accounting is a relatively new and technical area, and the accounting profession is only starting to address the important issues involved.

First of all, the hedge in question does not seem to meet the requirements for a hedge as stated in the GAAP. According to FAS 80, a futures contract must be marked to market at the end of the accounting period unless it qualifies as a hedge. To qualify as a hedge, (1) the futures must be designated by the firm as a hedge, (2) there must be underlying risk to hedge, (3) while the assessment of risk can be done on a centralized basis (if it is impractical to do otherwise), the risk management must be assessed on a decentralized basis for specific assets, liabilities, and firm commitments, and (4) there must be a clear economic relationship between the price of an underlying asset, liability, or firm commitment, and a high degree of price "correlation" must be *probable*. (The reference to correlation bears no relation to the usual statistical definition of correlation: FAS 80 makes it clear that the statistical definition is not intended and may not be relevant in assessing compliance. Unfortunately, FAS 80 does not make clear how correlation should be defined.) Under these rules, our hedge of sales certainly does not qualify, since future sales corresponding to use of copper in production are off balance sheet and are not firm commitments. Even if the sales were on the balance sheet, it is not clear whether they would meet the vague and mysterious requirement that correlation be probable.

Failure to qualify as a hedge often penalizes hedging. An unqualified hedge will typically reduce volatility of future cash flows but increase volatility of reported

earnings. This volatility is especially damaging when it causes violation of debt covenants or capital requirements imposed by regulators. Volatility of earnings may also subject management to criticism; given the current hysteria over derivatives, we may want to pardon a manager who forgoes an economically useful hedge to avoid the appearance of "risky exposure to derivatives." Part of the problem is that there seems to be no simple test, given the current state of hedge accounting, that the lay public can apply to distinguish risky speculation from good hedging.

One interesting feature of the accounting rules is that hedges that are economically equivalent may have very different accounting treatments. Suppose in the example above that demand does not depend on copper prices (putting the same number in all of the "Units Sold" column in Table 1A) and that we are simply interested in hedging the input cost at expected demand. Then it might seem equivalent to hedge through a long-term contract with a supplier, by buying copper futures, or by buying shares in a company whose share price tracks copper closely. However, the contract with a supplier has no impact on earnings before the actual sale, buying copper futures is covered by FAS 80 as discussed above, and shares in the copper company are accounted at fair value, but unrealized gains and losses are unlikely to appear in earnings (FAS 115). In each of these cases, there are various rules, ranging from somewhat specific (FAS 105 and 107) to incredibly vague (FAS 119), that require a company to report its risk exposure. FAS 119 is especially vague; basically, it calls on companies and accounting firms to come up with reports that can be used as a basis for later standards. This approach comes from a general recognition that current reporting practice is often misleading, and from a paucity of good ideas on how to patch things up. It seems that hedging tends to magnify the problems inherent in the accounting profession's tension between historical cost and mark-to-market cost.

It should be mentioned here that some people have proposed universal adoption of mark-to-market (or intrinsic value) accounting, which is “obviously” the correct thing to do because that is a good estimate of what the firm is actually worth, and any hedge would be seen for what it is. Unfortunately, it is not at all clear what this means. For example, do we include future sales in our valuation, and if so, how far in the future do we go, and how do we forecast and value the future flows? Anyone who has been involved in capital budgeting knows that estimates of future cash flows are often inaccurate and may reflect the forecaster’s optimism more than the prospects for the firm. Even without these conceptual problems, introducing a whole new system of accounting is not a trivial matter. While we note that current accounting standards are deficient for measuring risk, we do not claim that it is easy to do better.

The differences in accounting treatments of economically equivalent hedges may allow firms to hedge in spite of the deficiencies in the accounting standards. Whether or not a firm that is hedging properly can avoid looking bad, it is clear that a firm that is not hedging at all, or even increasing risk, can look fine.

COST ISSUES

What is the cost of hedging? It is tempting to think that the cost of the hedge is the cost of any securities purchased in the hedge program. In fact, the hedge is often bundled with an investment. It is a fair investment to buy a call option for its intrinsic value, and absent market imperfections there is no cost in doing so. In practice, the cost includes transaction costs such as commissions, bid-ask spread, and any internal costs of trading (e.g., hiring a trader and setting up accounting oversight). For publicly traded contracts in liquid markets, the costs are probably small and easy to measure. When hedging uses custom contracts provided by investment bankers, the costs are hard

to assess (because they are built into pricing) and may be much larger. On a more esoteric point, we may also want to include in the cost of hedging the alternative use of any capital tied up in the investment or in margin or variation accounts. On another subtle point, a hedge may be more costly than it appears if its pricing and tax treatment make it inappropriate for the firm.

What is the marginal cost that should be used as an input for decisions about pricing the output? It is probably common to use the hedged price, but in fact the marginal cost of the commodity at the time of use is the spot market price (assuming an active market that was probably necessary to implement the hedge in the first place). It is irrelevant that the price has been locked in for a fixed quantity, since that is sunk, and the profit will be collected or the loss borne on the hedged quantity however much or little is actually used. If more is needed, the shortfall will be purchased at the spot price. If less is needed, the excess will be sold at the spot price. In either case, the marginal cost is the spot price. If the marginal cost is taken to be the hedge price (or some average price), value may be discarded. For example, suppose the spot price is higher than the hedge price. Then a computation assuming that marginal cost equals average price or the hedged price would understate the true cost of buying more of the input, and additional units could be sold when it is more profitable to sell what can be produced from the hedged quantity of inputs.

What is the transfer price that should be used when the commodity is procured by one unit in the firm and used by another? For accounting purposes, the organization should decide up front how profits and losses in the hedging program will be shared. It is probably best to plan to do so in a way that hedges cash flows in each unit, since that will tie compensation in each unit more directly to performance within the manager’s influence. If sharing of hedge profit and loss is not decided in advance, an inherent unfairness may result. For example, suppose the transfer

price is ambiguous or renegotiable. If the transfer price is the market price when the market price is low but a hedged price when the market is high, the purchasing unit gets a “free option,” and the procuring unit loses—whether or not it is hedged. The free option allows the unit to buy at the hedged price or the market price, whichever is less. The procuring unit always loses money.

RISK-MANAGEMENT POLICY

Given that standard accounting procedures do not provide a particularly useful picture of the quality of a firm’s hedging program, it is especially important for management to adopt and implement an understandable and effective risk-management policy. Such a policy should specify the goal and scope of any hedging activity, and it should dictate the degree of centralization and the control systems. Furthermore, the policy should provide for oversight and evaluation of the effectiveness of hedging.

A common feature of the large publicized trading losses is a failure of control systems. Financial firms face a particular temptation to have inadequate controls. Because firms want to keep successful traders around, they may tend to be sympathetic to traders’ insistence that the bureaucracy should not interfere with their work. A failure to separate the operations and accounting functions from trading was an essential common thread in the recent losses of over a billion dollars each at Barings, Daiwa Bank, and Sumitomo. In each case, the loss was attributable to a single trader. It is important to devote serious talent to the job of monitoring traders, even though the monitoring job is less glamorous, somewhat unpleasant, and, when things are going well, seemingly unproductive.

Besides the scenario of speculation under the guise of risk management, risk management can be counterproductive if it is too localized. To illustrate, the example we discussed earlier showed

how a procurement department that is hedging material costs may actually make overall cash flows more variable if input prices tend to be high when the industry does well. Less damaging, but probably still wasteful, is the practice in which companies use different parts of the firm to offset hedging or they hedge economically irrelevant risks (such as risks that represent an insignificant part of a firm’s cash flow volatility). For most firms, the benefits of centralization (better control, economies of scale, and cost saving due to internal netting) will outweigh the costs (mostly the difficulty of communicating and aggregating needs). Of course, it is a good idea to have a formal policy in either case, whether risk management is centralized or dispersed.

A good risk management policy should state the goals of the hedging program. Is it the firm’s policy to hedge the value of the firm or, alternatively, earnings or dividends paid to shareholders, and if so, what risks should be hedged and what risks should be borne by the shareholders? Should hedging be implemented on a divisional or departmental level (to improve planning and incentive compensation) when that hedging does not reduce the overall variability of the firm’s value? Should the hedging program focus on cash flows, earnings, tax avoidance, or something else? We do not yet have definitive answers to these questions, but at least a consistent policy will minimize offsetting efforts.

One important (but probably often neglected) aspect of a risk-management program is the need for ex post evaluation. Especially because these programs are relatively new, it is entirely possible to design a program that is ineffective or that even increases risk (like the naive hedging strategy in our copper price hedging example). Only retrospective analysis of the results can verify that the program is actually reducing risk. The retrospective analysis should also look at any side effects of the hedging, for example variation or margin account payments required to maintain the hedge.

CONCLUSION

Risk management is an important and difficult area of corporate policy. We have seen news accounts of disastrous failures in risk management. Less spectacular, but perhaps more important, is the widespread use of futures contracts and swaps to hedge foreign exchange, interest rate, and commodity risks, since, without this ability to hedge, many profitable businesses would be too risky.

The next few years should be especially interesting, as companies work on implementing vague new accounting standards that require them to describe their risk exposure. Now is also an exciting time for the development of internal controls and policies as companies work on developing effective hedges while avoiding catastrophic losses.

SOME FURTHER READINGS

For general information on risk management, some banks have issued guides that may be useful. For example, the *J. P. Morgan/Arthur Anderson Guide to Corporate Risk Management* is a primer on risk management, while *The Chase Guide to Risk Management* is an extended glossary published by Chase Manhattan in association with *Risk* magazine. Another good general resource is "A Survey of Corporate Risk Management," which was a separately numbered insert to *The Economist*, February 10, 1996, pp. 1-22.

Some observers have debated the advisability of hedging with options. One criticism is that hedging with derivatives often amounts to gambling with firm money at the encouragement of banks ("Betting Your Hedges," J. Ralfe, *Risk*, July 1994, pp. 22-23). One good point is that customized option positions are unlikely to be a good value, especially if it might be necessary to unwind the position before maturity ("Caveat Emptor," D. Westby, *Risk*, June 1995, pp. 24-25). A different argument is that poorly constructed strategies with poor disclosure can lead to legal troubles for managers and directors ("Courting Trouble," W. Falloon, *Risk*, August 1994, pp. 32-33). On the other

side of the debate is the suggestion that hedging with options is a good idea and that most criticisms are unjustified ("Keep Those Options Open," M. Schewitz, *Risk*, October 1995, pp. 35-36).

There are, of course, many articles in the popular press citing specific hedging programs gone bad. The case of Metallgesellschaft has sparked a debate among academics, since it may be a case of miscalculation of the hedge. In particular, for commodities subject to temporary shortages, there is no reason to believe that price uncertainty over longer horizons can be hedged effectively by available traded options at short horizons, although this is often assumed to be the case. Several scholars have studied this question using theoretical tools (S. Ross, "Hedging Long Run Commitments: Exercises in Incomplete Market Pricing," 1995 mimeo, Yale University), empirical tools (F. Edwards and M. Cantor, "The Collapse of Metallgesellschaft: Unhedgeable Risks, Poor Hedging Strategy, or Just Bad Luck," *Journal of Applied Corporate Finance*, Spring 1995, pp. 86-105, or G. Bakshi, C. Cao, and Z. Chen, "Pricing and Hedging Long-Term Options," mimeo, University of Maryland), or both theoretical and empirical (S. Pirrong, "Metallgesellschaft: A Prudent Hedger Ruined, or A Wildcatter on NYMEX," mimeo, Washington University in Saint Louis). The defense of Metallgesellschaft's hedging program is that it was a textbook hedge that would have done fine if not interrupted (C. Culp and M. Miller, "Hedging a Flow of Commodity Deliveries with Futures: Lesson from Metallgesellschaft," *Derivatives Quarterly*, Fall 1994).

Value at risk ("An Overview of Value at Risk," D. Duffie and J. Pan, *Journal of Derivatives*, Spring 1997, pp. 7-49) is one methodology that is widely used in practice to quantify various common sources of financial risk. This methodology has its critics, both because an objective measure is difficult to agree upon ("VAR: Seductive but Dangerous," T. Beder, *Financial Analysts Journal*, September-October 1995, pp. 12-24), and because value at risk neglects idiosyncratic risk and some market sources of risk.