

MIT OpenCourseWare
<http://ocw.mit.edu>

15.997 Practice of Finance: Advanced Corporate Risk Management
Spring 2009

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.

Chapter 3: Why Companies Manage Risk

3.1 Which Side Are You On?

Although the many different parts of the firm all face problems involving risk, the problems are very different. As we mentioned in the last chapter, the problem facing a commodity trader is not the same problem that faces the business unit manager and not the same problem that faces the CFO. Because the problems facing each actor within the firm are different, they require a focus on different risk related issues.

It is helpful to think about the different offices and activities of a company along the same dividing line used to organize a typical accounting balance sheet: assets on the left and liabilities on the right. Figure 3.1 is a listing of different activities within a typical corporation organized in a layout that mimics a balance sheet's layout—assets on the left, liabilities on the right. Organizing the different activities according to this left/right, asset/liability split is instructive and useful.

Activities on the left are the traditional sources of value for non-financial corporations. Investments are made, operations are managed, and value is produced and realized through the sale of the company's products. This is where the firm's competitive advantage is usually thought to lie, whether in the low cost of its production, or the premium margins on its products, or the value of its R&D portfolio, or the organization of its human capital.

On the right-hand-side of Figure 3.1 are the auxiliary activities that support and facilitate the productive activities of the firm. These are the activities that are normally considered the domain of the financial office or related departments. This side includes the main interface with the capital markets.

Figure 3.1
Where Is Risk Management Relevant? A Balance Sheet Metaphor

Assets	Liabilities
Valuation of New Assets	Cash Management
Optimization of Operations	Tax Management
Product Pricing	Foreign Operations Reporting
Supply Chain Management	Interest Rate Risk Management
R&D Strategy	Accounting
Third Party Contracts	Debt Funding
Management Incentives	Equity Issuances
	Financial Flexibility

It is tempting to think of this right hand side as the natural locus of risk management activities. The buying and selling of derivative securities is often restricted to the treasury or other offices falling under the chief financial officer. When people think of “hedging” and “who hedges”, this is what usually comes to mind. And if risk management were purely a zero-sum game played on the financial markets, then this assignment would be true. But that’s a wrong, limited view of corporate risk management. It overlooks half of the ways that risk management raised shareholder value. Risk management happens on both sides of the balance sheet.

But the sources of value are different on the two sides.

The Left-Hand-Side: Taking Risks to Earn a Return

On the left-hand side are the various business units actually making investments, optimizing the operation of assets, choosing product lines and pricing structures, setting up supply chain structures, doing performance evaluation. For these activities, risk management is first and foremost a decision making tool.

If management is to build a successful business, it needs to be able to answer questions such as these:

- What is the value of that asset? How much is this extra risk worth?
- If I target a riskier value added product line, is the incremental profit worth the risk?
- What is the risk-reward tradeoff involved in delaying an investment, in developing a more flexible production scheme?
- Do I pay a premium to lock in a more reliable supply schedule? What is the downside risk worth?
- How long do I keep funding a speculative research and development program? When do I pull the plug?

In none of these situations is management trying to avoid risk. In none of these situations is the objective to hedge risk. Businesses make money by assuming risk. Every capital investment project involves spending cash up front in exchange for the gamble of future, uncertain cash flows. Expanding operations means expanding risk.

The task in all of these situations is to accurately assess, measure and price the risk, so that management can then make the right decisions about what to pay, how much to invest, how to operate assets, organize production, price and market products, and so on. Risk management is an input to all elements of good business management. The task is not to reduce risk, but to face it wisely.

In a few cases the correct business decision may be to lay off the risk, to pass it to another party.

- For example, this is done in the case of major construction projects by writing the terms of the contract to specify whether the contractor or the investor is responsible for cost overruns or changes in input prices. Assigning the risk of delays to the contractor is done as an incentive device, to motivate the contractor to perform on time, or as a negotiating device, to motivate the contractor to provide a realistic estimate up front. The investor does not have an aversion to the risks, per se—just an aversion to a contract structure that would create unnecessary risks.
- The Bombardier example from Chapter 2 is another case in which the company did not hold onto the risk created the money-back guarantee it offered to customers. It hedged this risk to its counterparty, Enron. In this case, the company is laying off the risk because it has no expertise in estimating the size and price of the risk. To make the right business decision on whether to offer the guarantee, it needs to know the cost, and the most reliable estimate of cost comes in the form of a simple price charged by an arms-length transaction with another party willing to accept the risk.

The profitable management of a business always involves choosing among alternatives, and may involve accepting some risks while avoiding others. What is the company's competitive advantage in risk taking? It should load up on those risks associated with its skills and competitive advantage. Other risks are unrelated to its skills and competitive advantage, and it will lay them off wherever possible to improve its focus and maximize the profitable use of its limited resources.

Risk management is about understanding the risks, the cost of risk, the price of risk, the value of risk, and knowing what choices are available, knowing where the firm can sell the risk, knowing what packages of risk can be easily sold into the market and what not.

On the left-hand-side, risk management is a problem of measurement, pricing and valuation. These risk management skills are all targeted to the problems of evaluating alternative investments, optimizing alternative operating strategies, comparing across alternative contract designs and terms, ranking alternative product pricing plans. Risk management enables the firm to value and price its activities more precisely, recognizing the varying quantities of risk in different activities, determining the market price for that risk, thereby better assessing the true profitability. Of course, with this power of more accurate valuation of risk comes the ability to compete better by offering a more finely tuned and targeted portfolio of products and services, the ability to lower costs by honing in on risky and therefore expensive activities and managing away the costly risk.

The value of risk management on the left-hand-side is direct: the company chooses more valuable assets, better prices its products, minimizes the risk-adjusted cost of delivering the same services, and so is more profitable.

The Right-Hand-Side: Economizing on Risks

On the right-hand-side, risk management is one of negotiating the firm's relationship with the capital markets. Here risk management is about reducing the company's cost of liquidity, increasing the firm's debt capacity, lowering the equity necessary to support a given line of business, lowering the firm's effective cost of capital and so expanding the resources available to it. Value is created on the right-hand-side because there are frictions in the capital market, and risk management greases the wheels, making the interface between the company and the financial markets operate more smoothly, more effectively, at less cost and with less disruption to the company's strategy.

The value of risk management on the right-hand-side is indirect. Risk management enables the firm to do more with less — more real assets with less financial capital. Starting from a limited pool of risk capital, risk management enables the firm to expand the scale of its operations on the left-hand-side. It is this expansion of profitable operations — the left-hand-side of the balance sheet — that generates value. But it is the risk management activities on the right-hand-side of the balance sheet, that indirectly makes expansion on the left-hand-side possible.

The indirect paths through which risk management generates value on the right-hand-side of the balance sheet are many, and the intermediate objectives risk management accomplishes are diverse. All of these intermediate objectives have as their end-point increasing shareholder value. In most cases (but not all) the risk management activities involve one form of hedging or another — i.e., reducing the company's exposure to risk. But beyond that, these different activities may have little in common.¹

Why does reducing risk raise the value of the company? In the next few sections, we'll briefly survey a couple of these ways.

Reducing the Cost of Working Capital

A large volume of what passes for risk management in corporations is targeted to little more than cutting down on the company's need for cash balances and for accessing the money markets.

The vast majority of hedging involves buying and selling futures contracts with a very short maturity—usually much less than one year. For example, a U.S. company may have invoiced a foreign buyer in its local currency, but wishes to lock in the dollar price. The company expects payment within 90 days and sells the currency forward, fixing the receivable in dollar terms. Many oil and natural gas field operators know the quantity they are likely to produce in the coming month or months, and they sell this amount forward, locking in the price at horizons of one, two or three months. This is called transaction hedging.

¹ A tragically good example in which risk management increased shareholder value by increasing risk is the case of savings and loan institutions (S&Ls) during the 1980s. Many of these financial institutions had experienced losses which threatened their long term viability. But because the government's loan guarantee covered those losses without charging the institutions a fee tied to the risk, it became advantageous for some to double up their risks. If things turned out well, they could continue as profitable operations. If things turned out poorly, the government was left covering the higher losses. Inevitably, many of the S&Ls went bankrupt, and the bill was enormous. This forced a major reform of the system and gave a useful push to more sophisticated risk-based capital regulation systems. See, for example, White, Lawrence, 1991, *The S&L Debacle: Public Policy Lessons for Bank and Thrift Regulation*.

The company's basic operating plans and decisions are unaffected by the hedging. The hedge is planned and executed after these basic operating decisions have been made. The hedge accommodates the operating decisions, not the other way around. Therefore transaction hedging does not *directly* improve the productive value of the firm's operations or investments.²

These short-run hedges also do not significantly impact the company's major exposure to exchange rates or oil and gas prices. The vast majority of the value of these companies is tied to the expected sales in years 2, 3, 4, 5 and beyond. If the exchange rate moves against the U.S. manufacturer, it will become less competitive and lose future sales, no matter how firmly it locks in the dollar value of sales already made. If the oil price collapses, the producer will drop in value based on the expectation of lost revenue on later sales.

So why hedge the near-term transactions and not the long-term ones? Why focus on sales already in the bag and ignore expected sales down the pike?

The answer is that the objective isn't really to lock in overall value, *per se*. The hedge is really meant to deal with a cash flow planning problem. The company's treasury has to anticipate the amount of liquidity it needs and keep on hand sufficient cash to meet its obligated payments. With enough notice, it can find the most efficient, low cost means to do so. But the dangers of variability in its daily cash flow obligations forces the firm's treasury to maintain a certain amount of its assets in liquid, low return instruments. The greater the variability in daily cash flow obligations, the higher the cash inventory it requires, and the greater the return forgone. If the company can lower the volatility of its cash flow obligations, it can reduce the cash inventory it holds and shift that capital into assets earning a higher return.

A second advantage of transaction hedging is that results across business units and across different quarters are made more comparable. The transaction hedge helps to lock in the return that management negotiated when it made the decision to execute the sale or to buy the product. Subsequently exchange rate and commodity price movements could wipe out the profit, or create a windfall that had nothing to do with management's decisions. Short-run hedging improves the measurement of certain types of performance.

² Another dubious motivation for transaction hedging is to lock in a favorable price in the forward market. While there may arise cases in which a corporate manager has some basis for believing that the forward price is better or worse than the expected future spot price, this is unlikely to be a reasonable foundation for shaping an ongoing transaction hedging policy.

Long run volatility in prices is central to the company's competitive decision and its major business decisions, but not to its short-term cash management strategy. In the short-run many key business decisions have already been taken, and the firm faces a simple exposure that can be readily hedged.

These short-term hedges do not in any significant way alter the firm's left-hand-side actions. But they do lower the cost of working capital that the company needs to keep in order to finance the activities on the left-hand-side of the balance sheet.

Liability Structure

Some forms of longer range hedging need not come in the form of trades in derivatives, but can be built into the regular liability structure of the firm. A company with significant revenues in a foreign currency may choose to issue a portion of its debt in that currency. This matching of the risk profile of the liability to the risk profile of the revenue stream increases the likelihood that the firm will be able to pay its debt even in the event of an unfavorable movement in the exchange rate. Minimizing the likelihood of a period of financial distress has the obvious benefit of saving on the direct costs incurred in such a crisis, costs such as lawyers fees, consultants fees, and myriad other burdensome charges. But the indirect costs of the threat of financial distress can be larger still.

The relationship between creditors and shareholders is smooth when financial distress is a remote possibility. As the threat of distress looms closer, the relationship becomes more conflictual and this has an impact on the firm's actual operations. Creditors begin to enforce loan covenants that restricts the firm's use of funds and its investments and asset disposals. Shareholders begin to concern themselves less with what maximizes the value of the firm as a whole, and more with what maximizes the long-run value of the stock, possibly at the expense of the firm and its creditors. These conflicts lower the value of the firm. Structuring the debt ex ante so as to minimize the probability that these conflicts arise, increases the value of the firm ex ante. Structuring the risk profile of the debt, tailoring it to the risk profile of the firm, is one tool for doing so. Risk management of the debt structure indirectly raises the value of the firm by reshaping the incentives of the creditors and shareholders so that the right operating and investment decisions get made in the future, even when the company's profitability is impaired. Changing the risk profile of the debt does not directly improve the stock value. Selecting a different risk profile does not get the firm a 'deal' on the price of its debt. But indirectly, through changes produced also on the left-hand-side, risk management on the right-hand-side is helpful.

Financial Policy and Strategic Management

Some corporations integrate risk management into their strategic plan. If the company's plans call for major capital expenditures, asset acquisitions or a well funded R&D program, then having the financial flexibility to complete these plans can be critical to shareholder value. Risk is no longer a simple gamble that sometimes yields gains and sometimes losses. Short-term losses may have large consequences. If the company is forced to cut back its capital expenditures or its R&D, then valuable opportunities are lost and the stock price suffers a double whammy. A well planned hedging strategy implemented up front can be used to assure that the company will have sufficient capital resources on hand to execute its strategy.

Once again, the value of the risk management strategy is indirect. The specific transactions used in the strategy – the gains and losses on a dynamic hedge portfolio, for example – are not the real source of value. But because the gains are made when the company needs cash, it is able to follow through on positive NPV investments and expand the value produced on the left-hand-side of the balance sheet. And hopefully the losses are made when the company has sufficient cash and therefore is not sacrificing extra value by forgoing negative NPV investments.

The strategic value of hedging comes not only in the direct form of being able to fund this or that planned investment, but also indirectly through how a company's financial flexibility affects its competition with industry rivals or its relations with key customers. Analyzing this strategic interaction can be very difficult, but also very important.³

3.2 The MM Theorem of Hedging

In addition to discussing how risk management adds to company value, it is useful to address one popular misconception about how hedging raising market value.

Students of corporate finance are familiar with the Modigliani-Miller irrelevance theorems. The essential core of the theorems is that the value of the firm is determined by the cash flow stream generated by its assets, i.e., on the left-hand-side of the balance sheet, and not by the amount of debt in its capital structure or the amount of the cash flow paid out in

³ Dasgupta, Sudipto, and Sheridan Titman, 1998, Pricing Strategy and Financial Policy, Review of Financial Studies 11, 705-737. See also Downie and Nosal (1998) who show a firm that risk management can help a firm with market power in the product market achieve a first-mover advantage over its rivals. See also Mello, Antonio and Martin Ruckes, 2005, The Role of Hedging in Product Market Rivalry, Working Paper.

dividends.⁴ Although the propositions were originally developed to address the choice between financing with debt or equity, and the choice of whether to pay a dividend or retain the cash, they are justly celebrated because the underlying principle behind the theorems has a very broad application to most all activities represented on the right-hand-side of the balance sheet. Hedging is one such activity.

The MM theorem of hedging says that, the value of the firm is independent of whether or not it hedges. The firm's value is determined by the cash flow stream generated by the assets, i.e., on the left-hand-side of the balance sheet. If the firm lowers the risk of its cash flow stream by selling a risky cash flow in the market in exchange for a low risk cash flow—i.e., by hedging—the value of the firm is unchanged.

Table 3.1 illustrates the MM theorem using as an example an oil producer valuing its next year's production. The first set of rows shows a valuation of the unhedged cash flow. The firm plans to produce 10 million barrels of oil. The forecasted spot price is \$35/barrel, so the firm's expected revenue is \$350 million. The risk-adjusted discount rate applied to this oil revenue stream is 10%, so the present value of this expected revenue is just over \$318 million.

The second block of rows shows the value of hedging this risky cash flow by selling futures contracts. We assume the hedge ratio (size of futures position divided by the size of the underlying quantity of the commodity being hedged) is 1, so that the company opens a short position in the futures contract for 10 million barrels. The futures price for this 1 year's production is \$33.09/barrel. The futures market for oil is deep and liquid, so we are entitled to think of this futures price as a competitive market price. We assume that the futures are settled financially and we abstract from issues of margining. Then at the maturity of the contract, the firm receives the difference between the futures price and the realized spot price. For valuation purposes, this is equivalent to the firm receiving the futures price and paying the spot price, so to value the hedge we value a cash inflow at the futures price and a cash outflow at the expected spot price.

⁴ Strictly speaking, absolute irrelevance is based on the assumption of no taxes. We make that assumption here as well, and address the problem of taxes later below.

Table 3.1
The MM Theorem of Hedging

Unhedged Valuation	
Forecasted Production (000 bbls)	10,000
Forecasted Spot Price (\$/bbl) -- current price \$38	\$35.00
Forecasted Spot Revenue (\$ 000)	\$350,000
Risk-adjusted Discount Rate, r	10.00%
PV (\$ 000)	\$318,182
Hedge Valuation	
Size of Position -- (000 bbls short, i.e., sold)	10,000
Futures Price	\$33.09
Forecasted Hedge Inflow @ Futures Price	\$330,909
Riskless Discount Rate, r_f	4.00%
PV (\$ 000)	\$318,182
Forecasted Hedge Outflow @ Expected Spot	\$350,000
Risk-adjusted Discount Rate, r	10.00%
PV (\$ 000)	\$318,182
PV Net Hedge Inflows & Outflows (\$000)	\$0
Net Position Valuation	
Total PV Forecasted Spot Sales + Hedge (\$000)	\$318,182

The expected cash inflow is \$330 million, while the expected cash outflow is \$350 million. So on average, we expect the short futures position to yield a negative cash flow. But it is not just the size of the futures cash flow and the expected spot cash flow that differ. The discount rates that are applied to the cash flows differ as well. Because the futures price is fixed at the time the hedge is purchased, we discount the expected cash inflow at the riskless rate. This is 4% in our example. Because the expected cash outflow fluctuates at the spot rate, we value the expected cash outflow at a risk adjusted rate reflecting spot price risk. This is 10% in our example. The present value of the two cash flows exactly match one another. They are both \$318 million. Therefore, the present value of the net position is zero. Adding the hedge on top of the firm does not change the firm's value.

One key feature of Table 3.1 is that present value of the cash inflow and the cash outflow on the hedge are equal. This must be true in an efficient capital market. The forward price,

discounted at the risk free rate, must be equal to the expected spot price discounted at an appropriate risk adjusted discount rate. If we change the forecasted spot price in the example, but keep the forward price fixed, then we must change the risk adjusted discount rate. If we do not make this change, then we are assuming a free lunch: anyone can buy or sell futures contracts at a profit.

Another underlying assumption is that there is an efficient capital market uniting both the valuation of the cash flow stream generated by the firm's assets and the cash flow streams on any hedge to be bought and sold by the firm. The price at which the high risk cash flow stream can be sold in the market is the same price at which the high risk cash flow stream earned by the firm would be valued by its shareholders. And that the price at which the low risk cash flow stream can be bought in the market is the same price at which the low risk cash flow stream would be valued by its shareholders. This assumption makes sense for a wide range of situations relevant to large corporations. For all intents and purposes, they sell their shares and their bonds in the same capital market that they can buy and sell hedges on their exposures. Of course, there are many cases where this assumption does not apply. For example, in any small firm where a single person has a majority of their own personal wealth tied up in the assets of the company.

The MM theorem of hedging does not, ultimately, say that hedging has no value to the firm. It just says the value does not arise *directly* from the exchange of a risky and riskless cash flow with different market values. The presumption is that hedging exchanges a risky and a riskless cash flow of equal value. So, in the first order, hedging leaves the value of the firm unchanged. However, changing the risk profile of the firm's cash flows can, *indirectly*, yield value to the firm. This indirect or second order benefit from hedging can arise in many ways as was described above and as will be explored in more detail in Part III of these lectures, The Business of Risk Management.

3.3 Risk-Aversion and the Value of Hedging

The MM theorem of hedging seems to fly in the face of a very large economics literature on the design of an optimal hedging strategy. Papers in that literature begin with a risk-averse agent who must buy or sell some commodity, and, depending upon the various assumptions about the risky price and other matters, they derive an optimal hedge ratio. In that literature, hedging directly benefits the agent. If the agent is a firm, then this directly contradicts the MM theorem.

How does one square these two approaches?

The MM theorems assume an efficient capital market. Given an underlying set of real economic activity with all of the attendant risks, these capital markets establish a market price for each type of risk. The total set of risks from the underlying economic activity is ultimately transferred to the global pool of investors. It is the risk-aversion of this pool of investors and their portfolio optimization choices that determine, in equilibrium, the market price for each type of risk. The firm will only be able to buy and sell risk in the capital markets at the equilibrium market price. The firm can hold more or less of any given risk, but it will pay the market price for this risk, leaving its total value unchanged. While the firm can hold more or less of any given risk via a hedge, ultimately the same total set of risk from the real activity of all firms still gets passed along to the global pool of investors. Hedging only changes the channels by which that happens. Whether a given risk is passed along as a part of the equity cash flow, or whether that risk is separated off and passed along as a part of a futures hedge, is irrelevant. The MM theorems assume that the firm's shareholders are just a random sample, indistinguishable from the rest of the global pool of investors, and that the shares they own of the firm are just a part of a well diversified portfolio. The same is true of the counterparties to the futures hedge. That's why the choice of channel doesn't matter and why the valuation of the risks is unchanged by any firm's hedging decisions. In the MM world, firms are just an intermediary shell where the real economic activity is organized. Firms pass along the value in the form of risky cash flows. The amount of risk channeled through the firm has no relevance for the value of the firm. The only thing that matters is the underlying set of economic activities. If those are unchanged – if the left-hand-side of the balance sheet is unchanged – then firm value is unchanged.

This is a very good assumption for a large segment of our corporate economy. And the MM theorem of hedging is a critical touchstone for students who want to understand risk management. But the assumptions behind the MM theorem don't apply to many other segments of the economy. When that is the case, then the firm's risk can directly matter in many important ways.

One such case in which the MM theorem of hedging does not apply is when shareholders are not well diversified. The founders of most start-up companies, for example, usually have a large fraction of their wealth tied up in the fortunes of that single company. Another example would be a nation that is heavily endowed with wealth from a single commodity such as oil. In these cases, these models with risk aversion make some sense, and there is obviously a potential value to hedging. But it is important to understand why this does not generalize to the standard publicly traded multinational corporation.

Models of risk aversion are also sometimes used as a reduced form representation of a much more complicated situation. For example, we may say that a company is risk averse within a given time frame and towards a specific risk because in a larger, more comprehensive model fitting the MM framework there arises naturally a premium price for this risk borne by the company. Managers of a company may make decisions for the company as if the company were risk averse because their own performance is evaluated and incentivized in a way that is equivalent to risk aversion. Again, it is important to understand the real source of this risk aversion.

In all of these cases, while a model of risk aversion makes sense and can be used to tell the moral of the story, it is very difficult to know how to parameterize the model and derive meaningfully calibrated quantitative results that can guide decision making. It is hard to know how to benchmark the measures of risk aversion. In the model of external capital markets, the price of risk can be benchmarked against market data on rates of return earned on various types of investment. Despite the many controversies over the right asset pricing model for measuring the price of risk, the range is narrowed by this data. The task is more difficult in these other cases. This is especially true when risk aversion is a proxy for indirect costs of risk, since the nature of the indirect costs varies on a case-by-case basis.

Earlier in this chapter we discussed the various right-hand-side of the balance sheet activities of the firm for which risk management may be useful indirectly. The indirect value in each of these cases is a little like the 'risk aversion as proxy' cases. There are particular internal activities and processes of the firm which can be executed more cheaply, efficiently or successfully given a carefully tailored risk profile for the firm's internal cash flows. These indirect benefits yield results much like what one would get by assuming risk aversion on the part of the firm. But as we model these right-hand-side activities in more detail later in these lecture notes, we derive the correct measure of the benefits. We do not assume a simple risk-averse utility function. In general, it is more insightful to model the actual factors driving the need for risk management and to let the model show us what sorts of risk management yield a benefit. It is seldom the case that the right answer matches what one would get from assuming a crude risk averse utility function.