

Enterprise Risk Management: Theory and Practice

July 2006

Brian W. Nocco and René M. Stulz*

* Chief Risk Officer, Nationwide Insurance, and Reese Chair in Banking and Monetary Economics, Ohio State University, NBER and ECGI. We are grateful for comments from Don Chew, Michael Hofmann, Joanne Lamm-Tenant and Jérôme Taillard.

Abstract

In this paper, we explain how enterprise risk management creates value for shareholders. In contrast to the existing finance literature, we emphasize the organizational benefits of risk management. We show how a firm should choose its risk appetite and measure risk when implementing enterprise risk management. We also provide an extensive guide to the implementation issues faced by firms that implement enterprise risk management.

There has been a dramatic change in the role of risk management in corporations. Twenty years ago, risk management often denoted the tasks associated with the purchase of insurance. Treasurers also performed risk management tasks, but they focused mostly on hedging interest rate and foreign exchange risks. Over the last ten years, corporations have taken into account additional types of risk. In particular, they started to pay much attention to operational risk and reputation risk. Most recently, strategic risks have been added to the panoply of risks considered. More and more, the risk management functions are directed by a senior executive with the title of chief risk officer (CRO) and the role of the board in monitoring risk measures and setting limits for these measures has increased at many corporations.

A corporation that chooses to manage risks can do so in two fundamentally different ways: it can manage one risk at a time, or it can manage all of its risks holistically. The latter approach is often called enterprise risk management (ERM). In this article, we argue that firms that succeed at ERM have a long-run competitive advantage over those that manage and monitor risks individually. Our argument is that, by measuring and managing its risks systematically and consistently and by aligning the incentives of employees to optimize the tradeoff between risk and return, a firm increases sharply the odds that it will be able to achieve its strategic goals.

In the following, we first explain why ERM creates value for shareholders and gives firms a competitive advantage. We then describe how ERM should be implemented. First, we explain how a firm should choose its risk appetite. Second, we show how it should measure risk. Third, we discuss the mechanisms that allow the firm to take and retain the risks that create value and lay off the others. Though ERM is conceptually straightforward, its implementation in practice is not. We therefore provide an extensive guide to the most important difficulties that arise in practice when implementing ERM in the last – and longest – section.

I. Why does ERM creates shareholder wealth?

ERM creates value both through its impact on the firm at a macro level and at a micro level. At the macro level, ERM creates value by enabling firms to quantify and rationalize the risk-return tradeoff they face and hence make it possible for them to access the resources over time to implement their strategy and to take risks that create value. At the micro level, ERM becomes a way of life for the corporation. All material risks are owned and the risk-return tradeoff associated with individual risks is internalized. Though the academic literature has focused on the macro level benefit for ERM, the micro level benefits are extremely important in practice.

A. The macro level benefits of risk management

Students in the first finance course of an MBA program often come away from such a course with the view that shareholders can diversify their portfolios, so that the value of a firm does not depend on its total risk.

Such a view of the world makes little sense.¹ Adverse cash flow outcomes have costs that go beyond the cash flow loss itself – in the language of economists, they have deadweight costs. If a firm expects a cash flow of \$200 million for the year and instead achieves a loss of \$50 million, a cash flow shortfall of this magnitude is generally more costly to the firm than the missing \$250 million. Obviously, such a cash flow shortfall has implications for the market's expectation of future growth, so that it might be associated with a reduction in firm value of much more than \$250 million for that reason alone. However, unless the firm was about to repurchase stock or debt for \$250 million or has excess cash of that amount or more, the cash flow shortfall means

¹ There is a large academic literature that investigates how firm value depends on total risk. See Stulz, R., 2002, Risk Management and Derivatives, Southwestern Publishing, for a review of that literature.

that the firm is short of funds and has to either cut back on its planned investments or raise additional funds. Raising additional funds is costly. Firms with debt capacity can issue debt, but raising debt to fund the cash flow shortfall reduces debt capacity and is often expensive. If the firm had what it felt was an optimal amount of leverage before the shortfall, it will have to raise a combination of equity and debt to maintain that optimal amount of leverage. The fact that so few firms issue equity in public markets, particularly in times of financial stress, suggests that for many firms issuing equity is too costly. When issuing equity has high costs, firms have little choice but to respond to the cash flow shortfall by cutting back investments that otherwise would be profitable. Firms forced to act that way leave money on the table. They give up what otherwise would be positive net present value projects and make it harder or impossible to achieve their strategic goals.

By managing risk, a firm can reduce the probability of large adverse cash flow shortfalls. Through hedging, firms can eliminate their exposure to many risks at low cost. For instance, a foreign exchange hedging program utilizing forward contracts has typically very low transaction costs. When hedging is cheap, there is no good economic reason for a firm not to hedge economic risks if it faces the possibility of cash flow shortfalls that could force it to give up valuable projects. Through hedging, the firm incurs small costs but increases the probability that it will be able to implement all positive net present value projects available to it.

Firms take many risks that they cannot profitably lay off in the capital markets or other developed risk transfer markets. For instance, a firm that finds a way to expand its business profitably cannot hedge cheaply many of the risks from doing so. The firm knows the risks from the expansion better than anybody else. If it seeks a way to hedge these risks, its hedging costs will be high because potential counterparties will want to be compensated for trading with a

better informed party and for constructing models to assess the risks they would hedge. It is not surprising that insurance companies do not offer insurance contracts that provide complete coverage for earnings shortfalls – the insured firm would know more about the distribution of its earnings than the insurance company and could manipulate that distribution to increase the value of the related insurance policy. Typically, a counterparty in a plain vanilla hedging transaction does not have to worry about dealing with a better informed party, so that for such transactions hedging is cheap. By taking on risks, firms make it more likely that they will have cash flow shortfalls that force them to reject valuable projects. Since firms cannot cheaply reduce their exposure to risks that are unique to their business, this makes it even more important for them to reduce their exposure to risks that are not unique to them. By doing so, they can take on more difficult to hedge, core business risks.

If there were no costs to hedging, firms would be best off hedging all risks. However, the costs of hedging risks increase as the risks hedged become more unique to the firm. A firm can cheaply hedge the exchange rate exposure due to foreign currency receivables. Hedging a Euro receivable is a commoditized transaction. However, it would be much more expensive for a firm to enter a contract that perfectly hedges the firm for exchange rate losses associated with sales that will be concluded over the next year because the level of such sales is uncertain and partly under the control of the firm.

B. The micro benefits of ERM

An increase in total risk is costly because it makes it more likely that the corporation will incur a shortfall that would force it to give up valuable projects. The cost of total risk therefore creates a risk-return tradeoff at the corporate level. If the corporation takes on a project that

increases the firm's total risk, that project has to be sufficiently profitable to pay for the cost of the increase in total risk. This risk-return tradeoff has to be evaluated for all decisions the corporation makes that have a material impact on the firm's total risk.

When implementing ERM, a firm has to ensure that decisions throughout the corporation will take into account the risk-return tradeoff correctly. Consequently, the firm has to make evaluation of the risk-return tradeoff pervasive throughout the corporation. It has to do so in three ways:

- a) First, the firm has to make sure that every material risk is taken into account in the measurement of total risk. If a manager undertakes a new project that has a material impact on risk, he will have to provide data to measure the marginal impact on the firm's total risk. Evaluating a project's return in relation to the marginal increase in firmwide risk will lead to optimization at the corporate level. Evaluating risk and return at a project level ignores diversification and correlations and will lead to suboptimal decisions.
- b) Second, a centralized evaluation of the risk-return tradeoff of individual projects would lead to corporate grid-lock. An extreme example is the case of a trader. A centralized evaluation of the risk-return tradeoff would involve the CRO in every decision the trader makes that has a material impact on the firm's risk. With a decentralized evaluation of the risk-return tradeoff, each unit in the corporation has to face this tradeoff in its decision making. The firm has to make it possible for each unit to resolve this tradeoff in the best interest of the shareholders.
- c) Third, performance evaluation of a unit has to take into account the contribution of the unit to the total risk of the firm. By doing so, the corporation not only measures the

true economic performance, but it also sets incentives for managers to manage the risk-return tradeoff effectively by not taking risks that are not economically attractive.

Through these three mechanisms that are essential to the management of firm-wide risk, a firm that implements ERM transforms its culture. Without these mechanisms, risk is taken into account suboptimally and/or subjectively when it is taken into account at all. In that case, some projects would be rejected because their risk would be considered too large even though the assessment may be wrong. Even worse, one division might take a project that another would reject based on their individual assessment of the cost of the project's risk. With the mechanisms put in place when ERM is implemented, the impact of all material risks on shareholder wealth is taken into account. Every risk is owned: it affects somebody's performance evaluation.

Risk ownership has become more important as the scope of risk management has increased to include operational and reputation risks. Early on, when risk management focused mostly on financial risks, a firm could centrally measure and change its exposure to market rates. However, operational risks typically cannot be hedged. Some of these risks can be insured, but often a firm can reduce its exposure to losses arising from such risks by changing its procedures and through mitigation. The individuals who are closest to the risks are in the best position to assess what steps should be taken to reduce the firm's exposure to these risks.

In practice, there are nuances relative to investment decisions and performance measurement which influence the implementation of a risk management and capital allocation framework. Nationwide Insurance employs a factor based capital allocation approach for its management accounting system. Capital factors are assigned to products based on risk, which provides business managers greater transparency with respect to the capital implications of their operating decisions. Factors are updated annually as part of the strategic and operational planning process,

reflecting changes in risk and diversification. Decision making is delegated through a risk limit structure consistent with Nationwide's risk appetite framework.

Significant investment decisions are analyzed at both the corporate and business unit level. Nationwide allocates diversification benefits within, but not across major business units. As such, Nationwide considers the impact on risk at the corporate level when making investment decisions while holding business managers accountable for earning adequate returns consistent with their competitive operating environment.

II. The right amount of risk

How should a corporation determine the optimal amount of total risk it should bear? Or, to use a language that is popular, what should the corporation's risk appetite be? To answer that question, it is important to realize that the adverse costs of cash flow shortfalls we discussed earlier would not exist if the firm had a larger buffer stock of equity capital invested in liquid assets. If the firm suffered a cash flow shortfall, it could then simply use its liquid assets instead. However, having such a buffer stock of equity capital is expensive. By reducing risk, a firm can lower the buffer stock of equity capital it requires. Consequently, firms face a tradeoff between equity capital and risk. ERM quantifies this tradeoff and enables firms to optimize it.

When the firm does not have an unlimited buffer stock of equity capital, a drop in the value of the firm can lead to a situation where the firm cannot pursue all the good projects available to it because it becomes financially constrained. Let's define financial distress to be a situation where a firm has to give up positive net present value (NPV) activities because it is financially constrained. There are different degrees of financial distress. For instance, when financial distress is moderate, one would expect a firm to be forced to give up only marginally profitable

projects, but not projects with large NPVs; as financial distress deepens, a firm may eventually reach a position where it cannot invest in new projects at all. Practically, however, differences in degree of financial distress are hard to assess. Instead, firms generally focus on a shortfall level that they want to control because exceeding that shortfall level has costs to the shareholders that go beyond the shortfall itself – i.e., if that shortfall level is \$100 million, a shortfall of \$150 million decreases shareholder wealth by more than \$150 million because it also costs the shareholders the loss of positive NPV projects. In general, a firm cannot guarantee that it will never have a shortfall that exceeds the threshold level it cares about. As long as a firm earns more than the risk-free rate, it has to take the risk of shortfalls. However, the firm can affect the probability of occurrence of such shortfalls. Minimizing that probability would not be in the interests of shareholders. The firm must trade off the probability of large shortfalls and the costs associated with such shortfalls with the expected gains from taking or retaining risks to optimize its risk portfolio.

Let's call the level of resources below which the firm is in financial distress the financial distress threshold. For simplicity, let's assume that this threshold corresponds to a level of capital below which the firm is in financial distress. It is often most intuitive for management to think of a bond rating below which the firm is in financial distress. Say that the firm incurs financial distress costs if its rating is Baa and below. Given a firm's current rating, it is straightforward, using data from the rating agencies, to estimate the average probability that it will have a rating of Baa or below. Suppose a firm has an Aa rating. A study from Moody's using data from 1920 to 1996 shows that firms with such a rating have a probability of 0.98% on average of having their ratings fall to Baa or lower over the coming year.² Such a probability may be too high for

² Moody's Rating Migration and Credit Quality Correlation, 1920-1996, July 1996.

some firms and too low for others. Consider a firm that has many valuable growth opportunities. For such a firm, an adverse shock that forces it to give up these growth opportunities will be very costly. Consequently, the expected loss to shareholders of financial distress will be large. A manufacturing firm with few growth opportunities may be best off maximizing the tax benefits of debt and returning funds to shareholders. For such a firm, financial distress would have low costs and it could afford a higher probability of financial distress. Financial firms have an additional consideration when they look at the costs of financial distress, which is that they make profits from some of their liabilities.³ For instance, bank deposits and insurance contracts are liabilities that result from core business activities of financial firms. These types of liabilities are very credit-sensitive, so that financial institutions typically maximize their value by targeting a much lower probability of distress than the typical industrial firm.

Suppose for the moment that a rating is a precise measure of the probability of default of a firm, so that if an economist were to compute the probability of default for two firms and were to find them to be the same, the two firms would have exactly the same rating. Consider a firm that would have to start giving up valuable projects if its rating fell to Baa or below. For such a firm, Baa is the rating threshold for financial distress. Let's assume that the firm has determined that the probability of financial distress that is optimal for its type of business is 6%.

As long as ratings are good proxies for financial health, the firm can theoretically use a rating agency transition matrix to estimate the amount of capital it would need for a given level of risk. The transition matrix identifies the frequency with which, over a historical sample period, firms moved from one rating to another. Table 1 reproduces the transition matrix estimated by

³ See Merton, Robert C, 1993, Operation and regulation in financial intermediation: A functional perspective, in *Operation and Regulation of Financial Markets*, edited by P. Englund. Stockholm: The Economic Council.

Moody's with data from 1920 to 1996.⁴ The column corresponding to a rating of Baa tells us the probability that a firm ends up with that rating at the end of the year for any rating at the beginning of the year. The firm wants the probability of its rating falling to Baa or lower to be 6%. To find out the probability of a rating equal to or lower than Baa for a given initial rating, we have to add up the probabilities of ending with a rating equal to or lower than Baa along the row corresponding to that rating. We added up these probabilities and they are shown in the last column to the right. We see there that the row with a probability closest to 6% is the row corresponding to a rating of A. In practice, this process can be far more complex. For example, Nationwide Insurance analyzes and manages its probability of default and probability of downgrade in separate but related frameworks. Probability of default is anchored to its target AA ratings and the corresponding bond default history. Probability of downgrade is influenced by managing risk concentrations, such as natural catastrophe and equity markets.

In this example, the firm maximizes its shareholders' wealth by targeting a rating of A. For a firm, equity capital provides a buffer that absorbs adverse shocks to enable the firm to avoid default. For a given firm, a different probability of default corresponds to each level of equity capital, so that if the firm chooses a given level of equity capital, it selects its probability of default. A firm with an A rating has a probability of default of 0.14% over a one year period. Consequently, to achieve that rating, the firm in our example must have enough capital that its probability of default is 0.14%. With default, all of the equity of the firm is consumed. So, the probability of default provides us with a way to determine the amount of equity the firm has to have given its current risk. The probability of default is typically a complicated function of firm characteristics, including the amount of equity. Intuitively, however, the process that leads to the

⁴ See footnote 2.

amount of capital is straightforward. Suppose that the firm becomes bankrupt if firm value at the end of the fiscal year falls below a threshold point K . The firm then needs an amount of equity capital so that the probability that its value will fall below K is 0.14%. If K corresponds simply to the principal amount of the firm's debt plus interest due, then the amount of equity capital the firm requires is an amount such that there is a 0.14% chance firm value will fall by more than this amount given the risk of its existing activities.

A firm can proceed beyond a focus on rating thresholds alone. For instance, Nationwide Insurance identifies numerous scenarios that it views as costly to the company. These may involve volatility of earnings and capital, or other key risk indicators, that individually may not materially impact ratings but effect overall risk and, therefore, the required level of capital. It sets target probability levels for these scenarios and identifies acceptable tolerances. It can then manage its volatility risk to maintain the probability of the various scenarios at their target probability levels.

Value-at-risk (VaR) is the loss that is exceeded at some pre-specified probability level over a defined time period. For instance, if a portfolio of securities has a one year VaR at the 5% probability level of \$20 million, there is a 5% chance the portfolio will have a loss exceeding \$20 million over the next year. When the determination of the buffer stock of equity proceeds along the lines described so far, the firm considered has to have an amount of equity equal to its firm-wide one year VaR determined at a probability level of 0.06%. VaR can be computed for a whole firm by assessing the distribution of firm value. For some firms, VaR has the same information as volatility, so that the firm can focus on the volatility of its value.⁵ However, for firms for which the distribution of firm value changes is not symmetric, the information provided

⁵ In particular, VaR is a multiple of volatility when the variable for which VaR is estimated has a normal distribution.

by VaR can be quite different from the information provided by volatility and for such firms VaR has to be estimated directly. For a given probability of default, the firm faces a tradeoff between its volatility or VaR and the size of its buffer stock of equity capital. Figure 2 shows this tradeoff. As the firm's volatility increases, it requires more capital to achieve the same probability of default. The tradeoff becomes steeper if the target probability of default falls.

Suppose that at the end of this process the firm finds that for its current firm value volatility it has to have equity capital of \$5 billion to achieve its target probability of default. The firm can reduce its required equity capital by reducing its risk of default. Reducing risk has costs. If the firm decreases risk through hedging, it faces hedging costs. If the firm decreases risk by eliminating projects that have high risk, the firm gives up the expected profits from these projects. However, equity capital has costs too. If the firm required less equity, it could increase its leverage and benefit from that increase. A decrease in equity capital might also increase shareholder wealth because investors would become less concerned about the possibility that the firm could use its equity capital to finance poor investments. At the margin, the firm has to be indifferent between decreasing risk and increasing capital. If it costs \$10 million to decrease risk by 1% and the firm saves \$10 million of equity capital costs by doing so, the firm does not increase shareholder wealth by decreasing risk and therefore has the optimal amount of risk.

With this approach to risk management, the firm can evaluate the marginal impact of a project on its risk of default and risk of financial distress. As total risk increases, the firm requires more capital to support that risk. The cost of the additional capital is a measure of the cost of the project's contribution to the firm's total risk. The project is worthwhile if it has a positive NPV before the cost of its contribution to the firm's total risk that is high enough to also cover that additional cost. Similarly, when evaluating the performance of a unit within the firm,

the unit contributes to shareholder wealth if its economic value added net of the cost of its contribution to the risk of the firm is positive.

It follows from this that the capital required to support the contribution of an activity to the risk of the firm is a measure of risk. It is a simple measure of risk that can easily be added up across activities or across risks. The strategic decision for the firm is the probability of financial distress that maximizes shareholder wealth. Conceptually, everything else follows from that decision.

We can now summarize how ERM works:

- 1) The firm has to select its risk appetite, which amounts to choosing the probability of financial distress that maximizes shareholder wealth. When ratings are used, the firm determines an optimal target rating based on its risk appetite and the cost of altering its probability of financial distress.
- 2) Given the firm's target rating, the firm estimates the amount of capital it requires to support the risk it takes. To do so, the firm estimates its risk of default.
- 3) The firm selects the optimal amount of capital and risk that yield its target rating. For a given amount of capital, the firm can alter its risk through hedging and project selection. Alternatively, for a given amount of total risk, the firm can alter its capital to achieve its target rating. At the margin, the firm is indifferent between changing its capital and changing its risk.
- 4) The firm decentralizes the risk-capital tradeoff, so that the whole organization makes decisions optimizing this tradeoff.

III. Implementing ERM

Conceptually, ERM is straightforward. However, implementing ERM is challenging. For a firm to be successful in implementing ERM, it is critical that the whole organization understand why ERM will create value. Employees have to understand that it is not an academic exercise but a critical tool to enable the firm to implement its strategy. Consequently, ERM has to be sold successfully at all levels of the organization. For the whole organization to espouse ERM, incentives have to be set accordingly. We consider in turn the main difficulties involved in making ERM work.

A. Inventory of risks

The first step in operationalizing ERM is to identify the risks the firm is exposed to. A common approach is to identify the types of risks that will be measured. Early on, financial institutions focused on market and credit risks. Eventually operational risk was added. As a result, a common typology for banks is to classify risks into market, credit, and operational risks. For such an approach to capture all the risks the firm is exposed to, operational risk has to include all the risks that are not market and credit risks. For banks, the definition of operational risk that prevails in the Basle 2 accord is much narrower – for instance, it ignores reputational risks. Consequently, there will be a tension between measurement of operational risk for regulatory purposes and measurement of operational risk from the perspective of ERM.

Many firms have gone beyond measuring market, credit, and operational risks. In particular, in recent years, firms have also attempted to measure liquidity, reputation, and strategic risks. Further, the typology used in banking often does not correspond well to the risks faced in other industries. In particular, insurance firms have risks on their asset side – their investment portfolio

– as well as on their liability side – the insurance payouts they may have to make. This leads insurance companies to use a different typology. For instance, Nationwide Insurance measures asset, liability, operating, liquidity and strategic risks, and considers reputational risks in the context of each of these risks and its overall business. The market and credit risks are subsumed in the asset risks.

If a firm follows the approach of classifying risks into market, credit, and operational risks, or uses a different typology that is more appropriate to its business, it then has to identify and measure how it is exposed to these various risks. This requires the identification and measurement of the exposures across the firm using common approaches. In the absence of common approaches, exposure to an identical risk would have different implications across the firm. Such a situation would create tensions when capital is allocated to various activities since identical activities could be allocated different amounts of capital. Further, eventually, risk would migrate within the organization to where its measurement is most lenient.

For an inventory of risks within an organization to be useful, it is necessary for all the information to be collected, made comparable, and updated. Organizations that have grown through acquisitions or without centralized IT departments typically face the problem that they have some computer systems that are not compatible. Firms need to be able to aggregate common risks across all of their businesses to effectively analyze and manage those risks.

Nationwide employs a top down and bottom up process of risk identification. From a top down perspective, its ERM leadership and corporate level risk committee have identified risks which are sufficiently large in aggregate that they could threaten financial distress in an adverse environment. The bottom up process involves individual business units and functional areas conducting risk control self assessments which are designed to identify all local level material

risks. The goal is to capture all risks, quantify risks employing a consistent approach and then aggregate individual risk exposures across the entire organization as well as analyze the aggregate risk profile considering risk correlations. For instance, Nationwide analyzes and establishes aggregate limits for equity risk represented by stock holdings in its property and casualty insurance investment portfolio, as well as fee levels tied to variable annuity and insurance contracts in its life insurance business and asset management fees in its investment management business.

Strikingly, inventory inadequacies are responsible for major corporate disasters over the last twenty years. It is often the case that business units push back against an inventory effort simply because such an effort is time-consuming and distracts from other activities. A well-known example of such push back that ultimately created massive problems for an organization is when one of the predecessor banks to UBS (itself called UBS) attempted to include its equity derivatives desk into its risk measurement. Because the equity derivatives desk used a different computer system, such a task required important changes in the way the desk did its business. The desk was highly profitable, so that it was allowed to stay outside the system. Eventually, it incurred massive losses that fundamentally weakened the bank and led it to seek a merger.⁶

B. Economic value versus accounting performance

Ratings are a useful device for a firm to determine its risk appetite. However, an actual rating is an estimate of a probability of default which depends on a number of accounting ratios as well as judgment from analysts. It is not therefore mechanically related to a firm's probability of default. As a result, a firm might position itself so that it has a probability of default

⁶ See Dirk Schütz, *La Chute de l'UBS, Bilan*, 1998.

corresponding to an A rating, but find itself in a position where rating agencies give it a Baa rating instead. If the firm uses the rating as a guidepost to assess its optimal probability of default, it might ignore the actual rating if it is comfortable that it has the appropriate probability of default. However, to the extent that the rating itself becomes a goal of the organization, setting capital at a level that achieves the probability of default of the targeted rating may not be enough. Most likely, the firm will have to target ratios that are important determinants of ratings as well. Such an approach could lead the firm to take actions that, while helping it to achieve the target rating, might lead the firm to focus more on accounting ratios than on economic value.

What is the shortfall the firm should be concerned about? Is it a cash flow shortfall? Is it a drop in a measure of firm value? Is it an earnings shortfall? Is it a shortfall in some key accounting ratio? If the corporation is managing its probability of default, it should obviously focus on the measure that makes it most likely that the probability of default is the one that it targets. An unexpected shortfall in this year's cash flow may not be a problem if future cash flows are not affected. If the corporation finds it easy to borrow against future cash flows or against tangible assets, a shortfall in this year's cash flow is unlikely to lead the firm to default. However, corporations that cannot borrow against future cash flows, perhaps because they are too speculative, and have few tangible assets, may be affected much more adversely by a shortfall in current cash flow. For such corporations, a shortfall in current cash flow can lead to financial distress because it makes them financially constrained. A firm in such a situation might choose to focus its risk management efforts on the distribution of its cash flow.

If the corporation is more likely to experience financial distress because the present value of future cash flows is low rather than because of a shortfall in current cash flow, it must model firm value and its distribution. Though top-down approaches that provide estimates of total risk

based on industry benchmarks are cheap and easy to implement, such approaches are not useful to manage risk within a corporation because they do not make it possible to relate corporate actions to firm-wide risk. For instance, a firm could obtain an estimate of the volatility of its value or cash flows by looking at the distribution of the value or of the cash flows of comparable firms. Such an approach would provide an estimate of the volatility of firm value or cash flows, but the firm would not know how specific actions affect this estimate. To implement ERM, the firm must therefore estimate the distribution of changes in firm value from the bottom up. Since firm value is the present value of the firm's future cash flows, the firm has to build its estimates by modeling the distribution of future cash flows. As part of its ERM, Nationwide Insurance has built stochastic models that generate multi-year cash flow distributions for its main businesses.

By focusing on cash flows, the firm focuses on its economic value. A firm may achieve a target probability of default and yet have volatile accounting earnings. Such an outcome can be even more likely under the current accounting treatment of derivatives. If a firm uses derivatives to hedge exposures and does not benefit from hedge accounting, the derivatives hedge can reduce the volatility of firm value but at the same time increase the volatility of accounting earnings since changes in the market value of the derivatives affect the firm's earnings. It follows that a firm that manages the volatility of the present value of its future cash flows or of its current cash flows does not necessarily manage the volatility of its earnings. A firm that implements ERM well could therefore have a higher volatility of earnings than one that does not.

Could it therefore make sense to focus instead on shortfalls in accounting earnings or to target a volatility of accounting earnings? It would for a firm which concludes that earnings shortfalls can have large costs even if not accompanied by commensurate cash flow shortfalls.

However, managing risk from the perspective of economic value accompanied by transparency around accounting earnings volatility likely leads to superior long-term results.

Nationwide Insurance operates in many businesses which are highly sensitive to credit ratings. As such, high levels of accounting volatility carry a risk to the extent ratings may be impacted. That notwithstanding, Nationwide attempts to make economically based decisions to maximize value creation. Ratings vulnerability is treated as a constraint. A significant amount of effort is devoted to minimizing the impact of this constraint through transparency and communication with the rating agencies.

C. Aggregating risks

A firm that uses the typology of market, credit, and operational risk generally measures each one of these risks. If the firm uses VaR, it will have a VaR for market risk, a VaR for credit risk, and a VaR for operational risk. It can then obtain a firm-wide VaR from these three VaRs.

The three types of risks have dramatically different distributions. This is also the case when risks are divided into asset risks, operational risks, and liability risks. Figure 2 shows how typically market, credit, and operational risks are distributed. Market risk behaves very much like the return on a portfolio of securities. The distribution of the return of a portfolio is typically symmetric. In contrast, both credit and operational risk have asymmetric distributions. With credit risk, either a creditor pays in full what is owed or he does not. In general, a majority of creditors pay in full, but some creditors default and when a creditor defaults the loss can be large. With operational risk, there are large numbers of small losses, so that small operational losses are almost predictable. There is also, however, some chance of some large losses, so that the distribution of operational losses has a long tail. Statisticians describe a distribution as having fat

tails if the probability of extreme losses is higher than can be described by the normal distribution. While many use the normal distribution to estimate the VaR of market risk, such an approach is not appropriate for credit and operational risks because these risks have fat tails.

When aggregating the risks, one must also estimate their correlations. The probability of experiencing simultaneously an extremely adverse market risk outcome, credit risk outcome, and operational risk outcome is typically very low. This means that there is diversification across risk categories. The firm-wide VaR is less than the sum of the market risk, credit risk, and operational risk VaRs. How much less depends on the correlation between these risks. The estimation of the correlations between certain types of risks is more art than science at this time and should consider the tendency for correlations to increase in highly stressed environments. It is often more common for firms to choose averages of correlations used in their industry rather than to rely on their own estimates.⁷

An important issue with correlations across types of risks is that these correlations depend to some extent on the actions of the firm. For instance, the total risk of an insurance company depends on the correlation between its asset risk and its liability risk. However, through changes in its asset allocations, the insurance company can modify the correlation between its asset risk and its liability risk. Consequently, an insurance company's asset portfolio allocations can be an essential part of its risk management effort.

⁷ For data on correlations used in practice, see Andrew Kuritzkes, Til Schuermann, and Scott M. Weiner, Risk measurement, risk management and capital adequacy in financial conglomerates, Wharton, Brookings-Wharton Papers on Financial Services, 2003, pp. 141-193.

D. Measuring risks

Some firms focus mostly on tail risk – the probability of adverse outcomes with very low probability. Consequently, when they measure the risk of changes in the present value of cash flows, they use a measure like VaR at a probability level corresponding to a default threshold. If firms complement VaR estimates with stress tests, where they investigate the impact on firm value of specific scenarios, they use as stress test scenarios rare events – such as the crisis period of August and September 1998 following Russia’s default on some of its debt.

Such a practice has been controversial and firms find that it reduces the usefulness of firm-wide risk management. One reason the practice has been controversial is that VaR measures a loss exceeded with some probability. It says nothing about the extent of loss if VaR is exceeded. Some have argued that firms should instead focus on the expected loss if VaR is exceeded. Focusing on such a risk measure, often called conditional VaR, instead of focusing on VaR has little economic justification in the context of firm-wide risk management. Setting the firm’s capital at a level equal to the conditional VaR would provide the firm with a lower probability of default than the one the firm targets since the conditional VaR exceeds the VaR.

A more important reason why it is important for firms to look beyond the VaR measure estimated at the probability level corresponding to a default threshold is that ERM creates value for corporations by optimizing the present value of financial distress costs. A firm typically has a much higher probability of financial distress than its probability of default. Further, as already discussed, different levels of financial distress have different costs, so that a firm can take these different costs into account and focus on the probability distribution of different levels of financial distress.

Estimating VaR becomes more of an art as the estimated VaR corresponds to a lower probability level. To see this, consider a firm that finds an A rating to be optimal. The probability of default for an A rating is 0.14% over a one year period. To estimate its optimal capital, the firm must therefore estimate the loss in value that is exceeded with a probability of 0.14%. Obviously, the firm has no experience of such a loss. Typically, it has no experience of a loss that comes close to that loss. Since the firm has little experience of such a loss, it is difficult for it to estimate the VaR at that probability level and to evaluate the outcome.

The firm can much better understand the distribution of changes in its value over the range of changes that includes rating downgrades rather than default. For instance, an A-rated firm has a 5.11% chance on average of being downgraded to a Baa rating over a one year period. Such an event is expected to happen more than one year out of twenty. In contrast, default is not expected to happen one year out of a thousand. The distribution of firm value changes corresponding to a downgrade to Baa may be well approximated using volatility. Over that range, the issues that arise because of asymmetries in the distribution of firm value changes and the so-called “fat tail” problems (i.e., the issue that extreme adverse outcomes are more likely than predicted by common statistical distributions) are not likely to be severe.

To the extent that the firm can estimate the distribution of value changes corresponding to a one grade downgrade much more precisely, managing the probability of downgrades should be an important dimension of firm-wide risk management.

As discussed previously, it is also important to understand and consider risk correlations when analyzing and managing default and distress probabilities. Nationwide Insurance incorporates a correlation matrix reflecting sensitivity tested stress correlations in its economic

capital model. It is also exploring event driven correlation analysis for scenarios including terrorist attacks, mega hurricanes and pandemics.

E. Regulatory versus economic capital

With ERM, the firm is concerned about creating value for its shareholders. It maximizes firm value by optimizing the tradeoff between risk and return. Adverse outcomes impose deadweight costs on the firm. The firm can avoid the adverse outcomes by taking on less risk and can reduce the deadweight costs by having a larger buffer of capital. With the approach discussed in this article, the buffer of capital that is relevant is capital in excess of the firm-value threshold corresponding to default.

In the simplest case, where the firm has one debt issue that matures at the end of the year and the firm has a one-year horizon, the debt repayment plus interest at the end of the year defines the default threshold since the firm would be in default if its value is less than that amount. In this case, the amount of capital required would be equity capital corresponding to a VaR with the appropriate probability level. The more realistic case is one where the firm has both short-term and long-term debt. In this case, firm value could fall below the sum of the face values of debt without a default occurring. With such a situation, the firm would have to have equity capital so that, given its current capital structure, the loss in value that would lead it to cross the default threshold has the target probability.

The amount of equity capital required for the firm to achieve its optimal rating may bear little relation to the amount of capital regulators require it to hold if the firm is regulated. A firm that practices ERM may therefore have an amount of capital that substantially exceeds its regulatory

requirements because it maximizes shareholder wealth by doing so. In this case, the regulatory requirements are not binding and would not affect the firm's decisions.

The firm could be in a more difficult situation if its required regulatory capital exceeds the amount of capital it should hold to maximize shareholder wealth. Nationwide Insurance refers to this as "stranded capital." To the extent that economic and regulatory capital are subject to different drivers, the difference between the two can be arbitrated to some degree to minimize the level of stranded capital. Nationwide allocates any residual stranded capital to its businesses and products. If all the potential competitors of the firm face the same onerous regulatory capital requirements, the capital the firm has to hold that is not justified on economic grounds is simply a regulatory tax. If some potential competitors could provide the firm's products without being subjected to the same regulatory capital, these potential competitors could offer the products at a lower price and the firm is at risk of losing its business to less regulated competitors. In this case, the firm would have to factor in the cost of regulatory capital of its various activities and would want to grow its portfolio of activities to consume less regulatory capital.

Regulatory capital for regulated firms is generally defined in terms of regulatory accounting. For ERM, firms focus on GAAP and economic capital. An exclusive focus on accounting capital is mistaken when accounting capital does not reflect correctly the buffer stock of equity available to the firm. The firm may have valuable assets that are not marked-to-market in its books but that it could sell or borrow against. With such assets, the firm's book equity capital understates the buffer stock available to it to avoid default. This suggests that the amount of GAAP equity capital available is only part of the story. The composition of the firm's assets matters as well. If the firm incurs a large loss and has no liquid assets it can use to finance it, the fact that it has a large buffer stock of book equity will not be very helpful. However, at this point, many firms

tend to evaluate separately their liquidity and the appropriate amount of equity capital they require. As the practice of ERM evolves, we would expect these firms to pay more attention to the relation between the optimal amount of equity and the liquidity of their assets.

F. Using economic capital to make decisions

If a firm could simply stockpile equity capital at no cost, there would be no deadweight costs of adverse outcomes. If an adverse outcome occurs, the firm could use its liquid assets to finance it and the adverse outcome would have no impact on the firm's investment policy. Firms are not in this situation. If the market perceives that they have excess equity capital, their value will be less because they are not earning their cost of capital on that excess equity capital.

When a firm undertakes a new risky activity, its probability of financial distress and its probability of default increase. Since financial distress is costly, this means that one of the costs of that new risky activity is its impact on the present value of the costs of financial distress of the firm. The firm can avoid these additional costs of financial distress by raising enough capital so that taking on the new risky activity has no impact on the firm's probability of financial distress. Consequently, the most straightforward way to estimate the cost of the impact of a new risky activity on the firm's total risk is to evaluate how much capital is required so that taking on that new risky activity has no impact on the firm's probability of financial distress. In practice, the increased capital required by the new risky activity could result from higher VaR at existing targeted ratings or by setting a higher target rating.

Suppose that before the firm takes on the new activity, the VaR estimate used to set the firm's capital is \$5 billion. Now, with the new activity, this VaR estimate increases to \$5.1 billion. Consequently, for the firm to have the same probability of financial distress as it had

before it undertook the new risky activity, it would need to raise capital of \$100 million. This capital would have to be invested in such a way that the investment does not increase the risk of the firm, since otherwise the VaR of the firm would increase further. If the risky new activity lasts one year and it would cost \$8 million to the firm to have this additional capital of \$100 million available for one year, the economic value added of the new activity should be reduced by this cost of \$8 million. If the firm ignores this cost, it effectively subsidizes the new risky activity since it ignores a cost of that activity. To the extent that riskier activities have higher expected payoffs before taking into account their contribution to the firm's probability of financial distress, a firm that ignores that contribution ends up favoring riskier projects over less risky ones.

Though the example just discussed is straightforward, the implementation of this idea in practice faces several difficulties. A firm is a collection of risky projects. At any time, a project's contribution to the firm's total risk depends on the risk of the other projects and their correlations. When business units are asked to make decisions taking into account the contribution of a project to firm-wide risk, they have to know how to evaluate that contribution. They cannot be told that the contribution will depend on everything else that is going to happen within the firm. As a result, when decentralizing the risk-return tradeoff, the firm has to enable business units to determine the capital that has to be allocated to a project to keep the risk of the firm constant with information easily available to them. Nationwide's factor based capital allocation system is an example of such an approach. This makes it necessary to focus on correlations of the project with the firm's known existing or projected activities. In this regard, Nationwide takes into consideration the extent to which each of its businesses is considered to be a core, long-term member of its operations when factoring in diversification benefits.

Often, firms are forced to simplify the problem even further and ignore correlations altogether when they set capital. In that case, the capital required to support a project would be set so that the project receives no benefit from diversification. The contribution of the project to firm-wide risk would then be considered to be the VaR of the project itself. To account for diversification benefits, the firm would then reduce the cost of equity. Similarly, when evaluating the performance of a business unit, the VaR of the business unit would be used to assess the contribution of the unit to the firm's risk.

G. The governance of ERM

How does a firm know that its ERM is performing well? While an outcome of effective ERM should be a better estimate of expected value and better understanding of unexpected losses, ERM does not eliminate risk. Consequently, adverse extreme outcomes still happen with ERM. ERM cannot, therefore, be judged on whether such outcomes happen. ERM manages the probability of such outcomes, but if the probability of default is set to one in a thousand years, there is no way to ever judge if the CRO did a good job in managing the risk of the firm so that the firm would have its target probability of default. Instead, to evaluate the job of a CRO, the board and the CEO have to evaluate how well the firm's risk is understood and managed. A firm where risk is well understood and well managed is a firm that can command the resources required to invest in the valuable projects available to it because it is trusted by investors. For such a firm, investors know that bad outcomes are the outcomes of bad luck rather than the results of a lack of control of management over the firm's activities.

IV. Conclusion

In this article, we have shown why enterprise risk management creates value for shareholders and examined the practical issues that arise in the implementation of enterprise risk management. The key principles that underlie the theory of enterprise risk management are well-established. However, it should be clear from this article that additional research is needed to help with the implementation of enterprise risk management. In particular, while much attention has been paid to measures of tail risk, such as VaR, it has become clear in implementations of enterprise risk management that a more complete understanding of the distribution of firm value is required. Though correlations between different types of risks are essential in measuring firm-wide risk, existing research provides little help in how to estimate these correlations when implementing enterprise risk management. Firms find some hard to quantify risks to be extremely important. Examples are reputation and strategic risks. At this point, there is little research that helps practitioners in assessing these risks, but much gain could be made by understanding these risks better even if they cannot be quantified reliably. Implementation of enterprise risk management has made a great deal of progress and shareholders have benefited. However, as this implementation improves with the help of academic research, we expect the gains from enterprise risk management to increase.

Table 1: Transition matrix from Moody's

Average one-year rating transition matrix, 1920-1996, conditional upon no rating withdrawal.

Source: Moody's Rating Migration and Credit Quality Correlation, 1920-1996, July 1996.

Rating From:	Rating To:	Aa	A	Baa	Ba	B	Caa-C	Default
Aaa	Aaa	92.18%	6.51%	1.04%	0.25%	0.02%	0.00%	0.00%
Aa	Aaa	1.29%	91.62%	6.11%	0.70%	0.18%	0.03%	0.07%
A	Aa	0.08%	2.50%	91.36%	5.11%	0.69%	0.11%	0.02%
Baa	A	0.04%	0.27%	4.22%	89.16%	5.25%	0.68%	0.07%
Ba	Baa	0.02%	0.09%	0.44%	5.11%	87.08%	5.57%	0.46%
B	Ba	0.00%	0.04%	0.14%	0.69%	6.52%	85.20%	3.54%
Caa-C	B	0.00%	0.02%	0.04%	0.37%	1.45%	6.00%	78.30%
	Default							13.81%

Figure 1: Required equity capital to achieve a target probability of default as a function of firm volatility or VaR

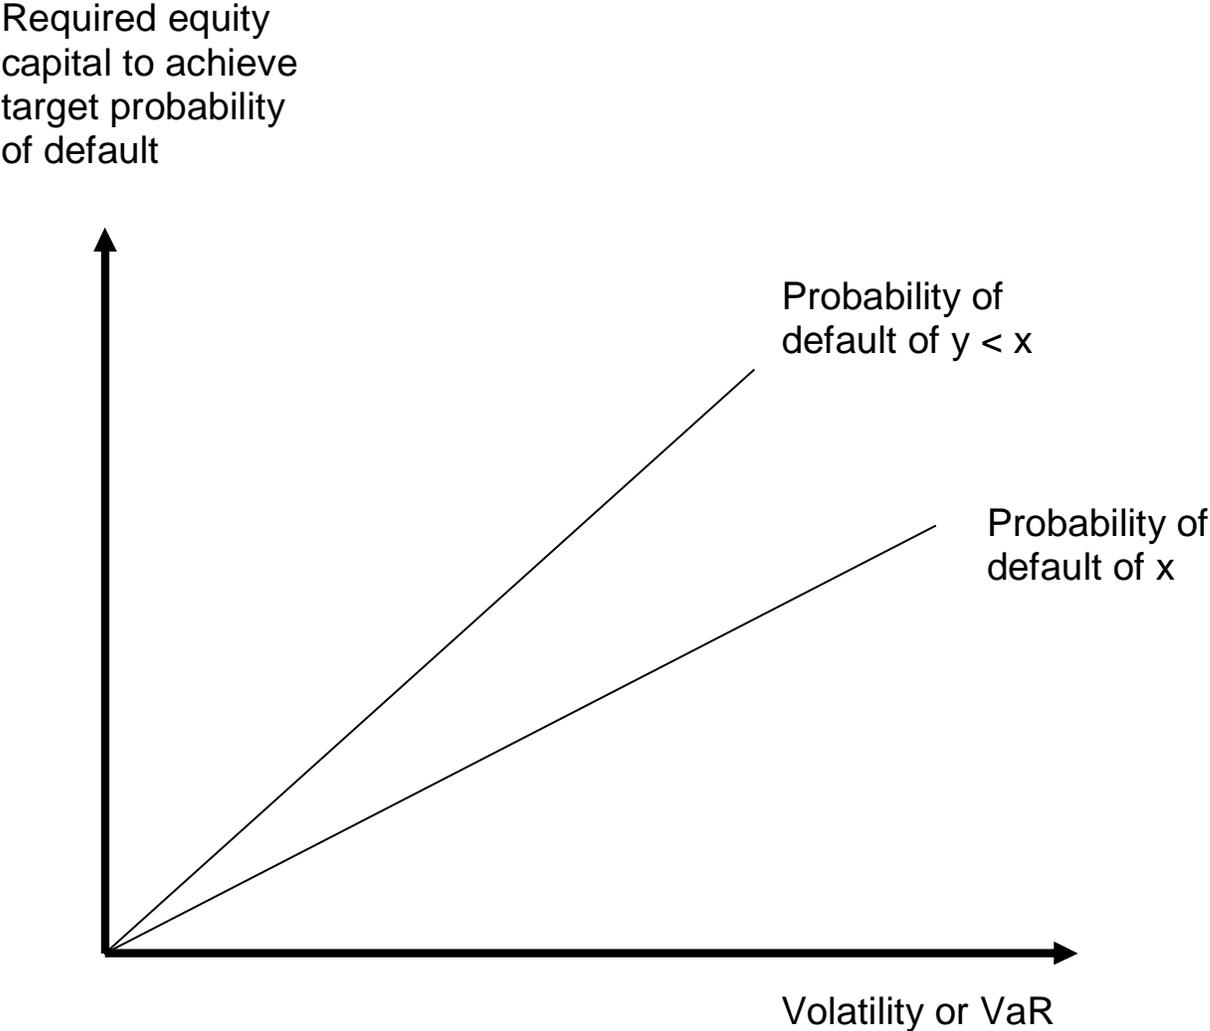


Figure 2: Typical market, credit and operational risk distributions

