

Information systems for risk management

by

Michael Gibson*

Federal Reserve Board

March 1997

Abstract

Risk management information systems are designed to overcome the problem of aggregating data across diverse trading units. The design of an information system depends on the risk measurement methodology that a firm chooses. Inherent in the design of both a risk management information system and a risk measurement methodology is a tradeoff between the accuracy of the resulting measures of risk and the burden of computing them. Technical progress will make this tradeoff more favorable over time, leading firms to implement more accurate methodologies, such as full revaluation of nonlinear positions. The current and likely future improvements in risk management information systems make feasible new ways of collecting aggregate data on firms' risk-taking activities.

* The author is a staff economist in the Division of International Finance, Board of Governors of the Federal Reserve System. I thank Brian Boyer and Garrett Ulosevich for excellent research assistance. This paper represents the views of the author and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, other members of its staff, or the Euro-currency Standing Committee.

1. Introduction

To lay a foundation for a discussion of the role of information systems in risk management, we must first define the business needs that drive financial firms to implement risk management functions. We see three such needs:

1. To better understand the risks it is taking, a firm wants to measure them. Risks that lend themselves to quantification, which are the only risks discussed in this paper, include market risk (the sensitivity of a firm's value to financial market variables like interest rates, exchange rates, volatilities, etc.) and credit risk (the sensitivity of a firm's value to default by its counterparties).¹
2. To provide better incentives to its business units and to individual employees, a firm wants to reward good risk-adjusted performance. The firm must measure its risk before it can adjust performance for risk.
3. To provide its shareholders with a consistent and optimal risk-return tradeoff over time, a firm wants to accurately match the amount of capital it employs with the risks it takes.

To meet these needs, firms have developed sophisticated risk measurement methodologies and have made substantial investments in risk management information systems.

Managers expect a risk management information system to provide them with the data they need to meet the above three business needs. Currently, most managers want four things from their risk management information system:

1. calculate Value at Risk;
2. perform scenario analyses;
3. measure current and future exposure to each counterparty;
4. do all three of the above at varying levels of aggregation, across various groupings of risks, across product types, and across subsets of counterparties.

With these four goals met, a manager can measure risk at the firmwide level, which is the level of aggregation that shareholders care about, and at the individual desk, product, or trader level, where decisions on risk positions and risk-adjusted compensation are taken.

This paper describes what a risk management information system must do to meet these four goals². In section 2, we use two simple examples to frame the problem as one of *aggregation*. We go

¹ Risks which are less easily quantified, such as legal risk (the sensitivity of a firm's value to legal judgements) and operational risk (the sensitivity of a firm's value to the performance of operational tasks), are outside the scope of this paper.

² Most large financial firms are moving toward firmwide risk measurement, but many are not there yet. Accordingly, some of the descriptions of risk management information systems in this paper apply not to current practice but to the information systems we expect to see 3 to 5 years hence.

into some detail of two approaches to calculating Value at Risk to see what is required of information systems under each approach. In section 3, we discuss several problems or choices that have to be addressed along the way. We highlight the interdependence between risk measurement approach and information system architecture in these two sections. In section 4, we discuss measuring aggregate market risk and the potential limitations that risk management information systems might impose. Section 5 concludes.

As an aside, we note that the risk management information systems processing requirements of a typical large financial trading firm are not huge compared with information systems that large firms in other industries have already implemented. Banks with large trading operations maintain large databases of transaction data, but phone company MCI has a three-terabyte customer database, growing at 100 gigabytes a month.³ The Frontier risk management system of Canadian bank CIBC gathers data from 160 organizational units worldwide daily to produce risk management reports, but retailer Wal-Mart, in its Arkansas headquarters, accesses weekend sales data from each of its 3,017 stores on Monday morning.⁴ Large trading banks maintain trading floors with dozens (or hundreds) of trader workstations, but each Wal-Mart store has dozens of cash registers, each of whose software is updated on average 24 times a year.⁵ Each trader may do dozens of trades a day, but Wal-Mart does 65 million transactions each week.⁶ However, financial trading firms may face an organizational (not technological) hurdle from their traditionally decentralized structure that makes it difficult to convince disparate trading units why a central risk management function needs their data.

2. Information systems requirements and risk measurement methodologies

The problem of designing an information system for risk management is a problem of *aggregation*. Data from each of a firm's trading locations worldwide must be aggregated to calculate Value at Risk or to perform a scenario analysis on the firm's worldwide portfolio. Different methodologies for calculating Value at Risk will require different slices of each trading unit's portfolio data to be aggregated across trading units, imposing different requirements on the risk management information system. In this section, we discuss some interactions between risk measurement methodology and information systems requirements.

³ "Towering Terabytes," Information Week, 30th September, 1996. (One terabyte = 1,000 gigabytes = 1 million magabytes).

⁴ "Wal-Mart Ups the Pace," Information Week, 9th December, 1996.

⁵ Ibid.

⁶ Ibid.

Table 1
Two simple examples

	Example 1	Example 2
Question to be answered	How many financial instruments does the firm have on its books?	How many counterparties does the firm currently have?
Information required from each trading unit	Number of financial instruments the trading unit has on its books	List of the trading unit's counterparties, identified uniformly across trading units
Calculation required to compute firmwide answer	Sum across trading units	Sum across trading units with duplicate counterparties removed

To clarify what we mean by aggregation, consider two simple examples, summarized in Table 1. In the first example, a firm wants to count the number of financial instruments in its firmwide portfolio. The centralized information system can take the number of financial instruments in each trading unit's portfolio and do a simple sum to get the firmwide answer. There are two ways the centralized information system can get the number of financial instruments in each trading unit's portfolio: either the trading unit can compute the number and provide it, or the centralized information system can use the trading unit's position data to do the count itself. Which way is chosen will depend on how flexible the trading unit's systems are and how easy it is for the centralized system to access the trading unit's data.

In the second example, the firm wants to count the number of counterparties to which it has current or potential future exposure. For this slightly more complicated question, the centralized information system needs a list of counterparties from each trading unit, and the counterparty identification scheme must be uniform across trading units. Again there are two ways the centralized information system can get the trading unit-level information: either the trading unit can provide it or the centralized information system can use the trading unit's position data to get the information itself. Again the choice will depend on the relative flexibility of the trading unit's systems and accessing its data.

In both examples, the centralized and decentralized approaches can give identical, correct answers to the question of interest. While we are reluctant to make absolute statements about how a particular firm would make these choices, we are comfortable making two relative claims about the tradeoff between the flexibility of trading units' systems and the ability of a central risk management function to access position data.

1. Comparing a simple query (Example 1) with a more complicated query (Example 2), the second example's more specific information needs make a centralized solution more likely than in the first example. Producing a list of counterparties with a uniform counterparty identification scheme requires more flexibility from the trading unit's systems than simply counting the number of financial instruments on the books.

2. A centralized solution can handle different queries with the same data, while a decentralized solution requires each trading unit to provide different data for different queries. In other words, a centralized solution is more open-ended.

2.1 Aggregation and Value at Risk

Computing Value at Risk on a firmwide basis is also an exercise in aggregation, though more complicated than the simple examples presented above. The data on portfolio composition that the risk management information system needs to access and the types of calculations it has to perform will depend on the methodology chosen to calculate Value at Risk.

Many methodologies exist to calculate Value at Risk, defined as a portfolio's maximum loss over a given time period with a given probability. Each methodology combines an assumption on the future distribution of market risk factors and current data on portfolio positions to approximate the distribution of the change in portfolio value. Certain methodologies can severely test the ability of the risk management information system to gather the necessary data and do the necessary calculations in the time frame required. Consequently, firms often trade off accuracy and computational demands when computing Value at Risk.⁷ We discuss two Value at Risk methodologies: delta-normal and full revaluation Monte Carlo.

2.2 Delta-normal

The delta-normal methodology (also called the J.P. Morgan RiskMetrics™ methodology) stipulates that (1) the future distribution of changes in market risk factors is assumed to be multivariate normal; (2) portfolio positions are summarized by the "deltas" of each position with respect to each market risk factor, where "delta" is the change in the position's value for a one-unit change in the market risk factor; and (3) the distribution of the change in portfolio value is approximated linearly by the sum of the products of the firmwide delta for each market risk factor and the assumed distribution of that market risk factor.

Any Value at Risk methodology must make some assumption on the future distribution of changes in market risk factors. Possible assumptions include multivariate normality, another multivariate distribution, or using the historical distribution of changes to proxy for the future distribution. If a parametric distribution is used, the parameters must be estimated from some combination of historical data and current data such as futures, forwards and options. One important burden on the risk management information system is to maintain a database of historical time series on the relevant market risk factors to be used to estimate the covariance matrix of future changes.

⁷ See Pritsker (1997) for an evaluation of the tradeoff.

Because this burden does not vary significantly with the choice of methodology, and because it can easily be contracted out,⁸ it will not be discussed in what follows.

For each trading unit, "deltas" must be computed for each market risk factor. The RiskMetrics™ way to compute these deltas is to decompose each instrument into a sum of positions in some subset of the market risk factors. (This is equivalent to taking a linear approximation or a first-order Taylor expansion.) For example, if the set of market risk factors includes zero coupon bonds but not coupon bonds, a coupon bond would be decomposed into a set of zero coupon bonds of different maturities. An option on a stock market index would be decomposed into a position in the index itself. Deltas for each instrument in the trading unit's portfolio are summed to give the trading unit's deltas; these are summed across trading units to yield the portfolio's deltas. One advantage of the delta-normal method is that deltas can be easily aggregated. Finally, estimating Value at Risk requires taking the square root of a weighted sum of variances and covariances of market risk factors, where the weights are simple functions of the portfolio's deltas.

Calculating a trading unit's "deltas" requires access to the cash flows of each instrument in the trading unit's portfolio. As in the two simplified examples given above, there are two ways this can be done. If the trading unit calculates its own deltas, it must do so using a uniform set of market risk factors so the results can be aggregated at the firmwide level. For a system with N market risk factors this would involve computing and passing N "deltas" for each trading unit. However, any task that involves specifying new computations to be performed by each trading unit's systems will be more difficult to carry out the more diverse those systems are. If the centralized risk management information system calculates the deltas, accessing the portfolio data of each trading unit becomes a burdensome requirement of the delta-normal method; the difficulties of accessing transaction data stored in disparate systems often lead firms to duplicate their transaction data and store it centrally. While the location of the computation would differ, the estimated Value at Risk would be unaffected by the choice of centralized or decentralized processing.

The normality assumption and the delta (i.e., linear) approximation together imply that the distribution of the change in portfolio value is easy to compute. It will be normally distributed, and its variance will be a weighted sum of the variances and covariances of the market risk factors on which its value depends, with the weights depending on the "deltas" of the portfolio's positions. Since the probability percentiles of the normal distribution are well known, once the variance of the change in portfolio value is known the Value at Risk can be computed immediately.

To summarize, the delta-normal method requires that the risk management information system know the "deltas" of each trading unit's portfolio, which requires access to a large amount of

⁸ Currently J.P. Morgan does this for free, providing a covariance matrix each day on its Internet site (<http://www.jpmorgan.com/RiskMeasurement/RiskMetrics/RiskMetrics.html>). Olsen and Associates also provides a covariance matrix via the Internet (<http://www.olsen.ch/cgi-bin/w3risk-menu>).

data - the cash flows of each instrument in the portfolio. The computational burden is relatively light, requiring a large number of simple calculations to be performed.

2.3 Full revaluation Monte Carlo

When using the full revaluation Monte Carlo methodology to calculate Value at Risk, the future distribution of market risk factors need not be assumed to be multivariate normal. Some parametric distribution must be chosen, and its parameters estimated, but because normality is not required in this methodology an alternate distribution that better captures the statistical features of financial time series can be used.⁹ Two examples of such distributions are Student's t distribution and a mixture of multivariate normal distributions. The distribution of changes in portfolio value is approximated by taking a large number of draws from the assumed distribution of changes in market risk factors and revaluing the portfolio for each draw. For N draws, the 5 percent Value at Risk would be the $(.05N)$ th largest loss.¹⁰ Value at Risk can be computed in the same way, using the same N draws, for an individual trader, individual desk, or firmwide.¹¹

Each trading unit's portfolio must be revalued for each of the N Monte Carlo draws. These revaluations will be computationally burdensome for some instruments. In particular, complex derivatives that cannot be valued analytically are typically solved by computation-intensive numerical methods. Valuing such a derivative for, say, $N=10,000$ Monte Carlo draws would require a significant amount of computing power, much more than the linear approximation of the delta-normal methodology. On the other hand, two factors mitigate the computational burden of full revaluation Monte Carlo and one factor makes it easier to deal with. First, firms have an important business need to speed up the valuation of complex derivatives, for trading purposes as well as for risk management. There is no reason to think that research into faster numerical option pricing methods will not continue to be fruitful.¹² Second, as discussed in section 3.4 below, "smart" valuation techniques, such as using an analytic approximation to value a complex option, could significantly reduce computational burden with only a small or no reduction in accuracy. Finally, parallel processing techniques will not reduce the computational burden of multiple portfolio revaluations, but such a burden would be relatively easy to divide across many processors and thus reduce computational time at the expense of additional computer hardware.

⁹ The high dimensionality of the set of market risk factors rules out a nonparametric approach in nearly all circumstances. An alternative methodology, not discussed in this paper, takes a nonparametric approach by using historical data to represent the future distribution of market risk factors.

¹⁰ Pritsker (1997) shows that an advantage of full revaluation Monte Carlo over delta-normal is that a confidence interval on the estimate of Value at Risk can be computed at no additional computational cost.

¹¹ This may not be the most useful way to measure the marginal market risk of a trading unit, since it does not account for diversification across trading units.

¹² To give one example of such research, Carverhill and Clewlow (1994) describe how to speed up Monte Carlo valuation of options by a factor of 70 with a martingale variance reduction technique.

The calculations needed to revalue the trading unit's portfolio can be done either at the trading unit or at the central risk management function. If the trading unit revalues its portfolio for each Monte Carlo draw, the only burden on the central risk management function is to ensure that each trading unit uses the same N Monte Carlo draws to revalue its portfolio. Each trading unit would then simply pass N numbers, representing the changes in the value of its portfolio under the N Monte Carlo draws, to the central risk management function for aggregation. If the central risk management function does the revaluation of each trading unit's portfolio, it will again require complete position data from each trading unit along with a valuation model for each instrument. This last requirement is nontrivial for some complex derivatives, for which no "market standard" valuation technique exists.¹³ Again, the choice between centralized and decentralized processing need not affect the estimate of Value at Risk.

Once the change in portfolio value for each of the N Monte Carlo draws has been calculated, these N changes in portfolio value will approximate the distribution of the change in portfolio value and can be treated as an empirical distribution function. The Value at Risk at confidence level α can be read off the ordered list of N changes in portfolio value as the αN^{th} largest loss.

The full revaluation Monte Carlo methodology makes greater demands on the risk management information system than the delta-normal. Randomly drawing a large number of changes in market risk factors from a multivariate distribution can be done easily and quickly,¹⁴ but computing the change in portfolio value by revaluing each trading unit's portfolio for each draw from the assumed distribution of changes in market risk factors will be computationally burdensome for some instruments.

2.4 Choosing a risk measurement methodology

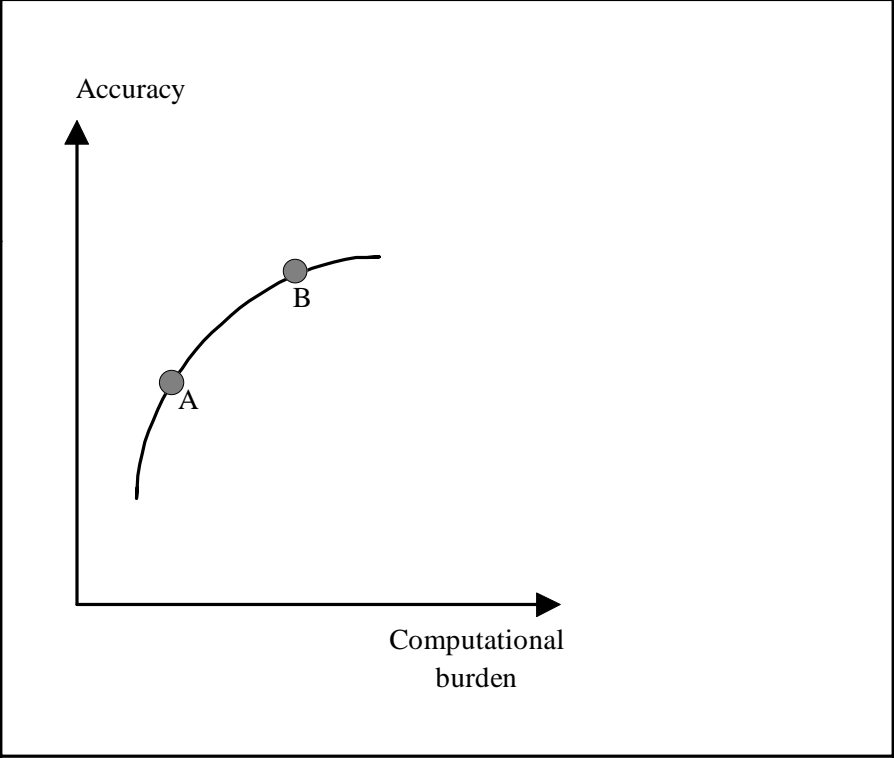
In choosing between a delta-normal or full revaluation Monte Carlo methodology to measure Value at Risk, a firm will trade off the accuracy of its Value at Risk estimates with the computational burden required to compute the estimates. The tradeoff can be represented as a curve, as shown in Figure 1. The delta-normal method would be a point like A, with relatively low accuracy and low computational burden. The full revaluation Monte Carlo methodology would be a point like

¹³ See Bernardo and Cornell (1997) for an example of an auction of mortgage-backed securities and the diverse valuations of large broker dealers and institutional investors. Pierides (1996) shows that the price of interest rate derivatives can be sensitive to the stochastic process chosen to model the short-term interest rate.

¹⁴ Drawing large numbers of random vectors from a high-dimensional space presents its own computational problems. Press et al (1992, p. 277) point out that typical computer random number generators cannot fill up a high-dimensional space. A discussion of random number generation is beyond the scope of this paper, but we conjecture that drawing truly random points in a high-dimensional space will require some thought at the time of system design but not a significantly increased number of calculations.

B, with relatively high accuracy and high computational burden. Each firm's choice will depend on the relative importance of the two factors for that firm.¹⁵

Figure 1



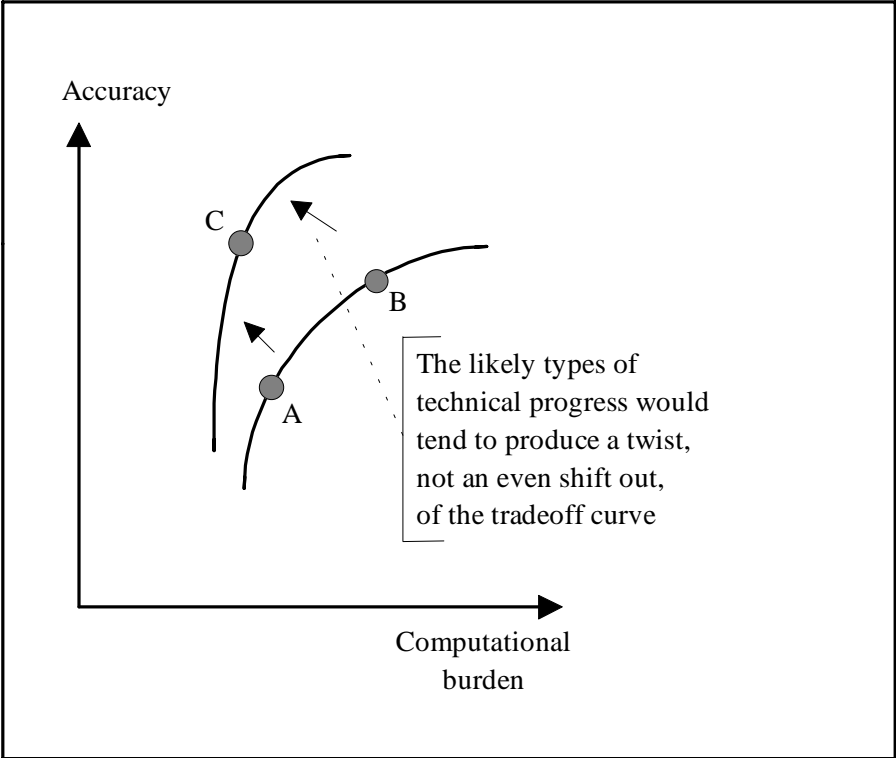
While we cannot predict where an individual firm will choose to be on the tradeoff curve in Figure 1, we do feel comfortable predicting how the curve will shift over time. Technical progress - including faster option price techniques, cheaper computer hardware, and advances in computer networking - will cause the curve to shift out, making the tradeoff more favorable. Such a shift is shown in Figure 2.

In addition, we can predict what the shift of the tradeoff curve will look like, as Figure 2 shows. Financial theorists will likely continue to produce faster option pricing models, and faster computer hardware will make a large number of Monte Carlo draws less burdensome to handle. Because the delta-normal methodology is already quite simplified, neither of these advances would reduce its computational burden by much. On the other hand, the scope for reductions in the computational burden of the full revaluation Monte Carlo methodology from these advances is large. For this reason, we predict that the reduction in computational burden, holding accuracy constant, will be greatest for those methodologies that achieve high accuracy, producing a twist in the tradeoff curve as shown in Figure 2. As this twist occurs over time, firms are likely to switch away from a

¹⁵ We know that this tradeoff captures a meaningful choice, since there are firms using each of the two methodologies.

low-accuracy delta-normal methodology to a high-accuracy full revaluation Monte Carlo methodology.

Figure 2



3. Other issues to be addressed

Our discussion of the information systems requirements of the two Value at Risk methodologies has touched on some interactions between methodology and information systems design. In this section, we discuss some problems or choices that any firm setting up a risk management information system must confront.

3.1 Centralized or decentralized?

A basic decision that must be made when designing a risk management information system is whether to choose a centralized model or decentralized model. As the discussion above of two simple examples and two Value at Risk methodologies should make clear, either model can estimate Value at Risk or do scenario analysis. The two models will give identical answers if given identical data. The difference lies in where the most burdensome calculations are done. Which model is chosen will depend on several factors: the flexibility of trading units' systems, the degree of uniformity among trading units' systems, the availability of a large central database to hold position data, and the

ability of the central risk management function to handle and revalue all instruments traded by all trading units.

The centralized model has several strong points. The risk management function can more easily monitor that the risk calculations (computing "deltas" or doing revaluations) are being done properly if they are being done centrally. In particular, if risk measurement calculations are being used in risk-adjusted compensation calculations, the centralized model avoids moral hazard. It is easier to "upgrade" to a more rigorous risk measurement methodology (from delta-normal to full revaluation Monte Carlo, for example) if a centralized model has already been adopted so that position data is already being accessed centrally. Finally, the centralized model is more open-ended, and a firm may find other uses for a centralized database of position information beyond risk measurement.

A decentralized model also has advantages. There is no need for position data and the analytics needed for revaluation to be duplicated at the central risk management function. The analytics for complex derivatives may be particularly costly to duplicate. Because there is no need to access position data or do revaluation in the central risk management function, there is no obstacle to introducing new products. For example, a data architecture that provides one field to identify the underlying security for an option would have a hard time handling a compound option whose value depends on two or more underlyings. If the analytics behind a compound option are new and complicated, it may be costly to duplicate them at the trading unit *and* central risk management level.

3.2 Mapping

Choosing a set of market risk factors on which to base a risk measurement methodology is an important decision. A firm making the common choice to use the RiskMetrics™ set of market risk factors or the default set of market risk factors in the risk management software it has purchased is likely not recognizing the importance of this decision. The choice of the set of market risk factors is equivalent to defining the "market risk" to be measured; risk left unmeasured, which can be termed "basis risk," is not guaranteed to be small and depends on how the set of market risk factors is chosen.

Consider a firm that has some exposure to the NLG/USD and DEM/USD exchange rates. If the firm's set of market risk factors includes both exchange rates, it will be able to measure its exposure to all possible combinations of moves in the two exchange rates, including moves in opposite directions. If the firm's set of market risk factors includes only the DEM/USD exchange rate and NLG/USD exchange rate risks are mapped into DEM/USD exposures for risk measurement purposes using the historical correlation between the two exchange rates, no market risk measurement technique will be able to reflect the possibility that the correlation could (with low probability) change. That possibility will fall under basis risk, not market risk, in the second case. Delta-normal Value at Risk will still be measured correctly, since with that methodology the correlation is held to be constant over the time period during which Value at Risk is calculated. However, the output of both stress tests and sensitivity analysis with respect to individual market risk factors, two commonly

used supplements to Value at Risk, will depend on the mapping. Two firms with identical positions but different sets of market risk factors will get different stress test or sensitivity outcomes.

A firm chooses the set of market risk factors on which to base its risk management methodology with several considerations in mind. If fewer market risk factors are chosen, the methodology will be easier to work with, to understand, and to explain, and fewer computing resources will be needed. If more market risk factors are chosen, the approximation error involved in mapping a position onto a limited number of market risk factors can be reduced.¹⁶ If more market risk factors are chosen, (unmeasured) basis risk can be reduced and (measured) market risk expanded. A firm will be able to minimize basis risk if it includes in its set of market risk factors every market risk factor to which it has exposure. Of course, exposures change over time, while it may be costly to update information systems to expand the set of market risk factors, so it is unlikely that basis risk can be eliminated as a concern.

3.3 Legacy systems

Another problem to be addressed when building a risk management information system is that of legacy systems. The term "legacy system" here refers to a trading unit's information system that cannot be easily integrated with a central risk management system. Although the typical legacy system is an old system that lacks features that have only recently been invented, the problem of legacy systems can include newer software as well, if that software cannot easily be set up to interface with a central risk management system. For example, if the pricing model for a complex option exists only in a trader's spreadsheet, it may be impossible for the central risk management system to request that the option be repriced for N Monte Carlo draws, as a full revaluation Monte Carlo Value at Risk calculation might require.

While the problems of legacy systems are serious, they should not significantly hamper a firm's ability to do firmwide risk management. Because legacy systems are so widespread, not just in financial services but in all industries, solutions to legacy systems problems are common. Many consulting firms exist solely to provide these solutions. One common solution is to write a "wrapper" program that acts as a mediator between the central risk management function and the trading unit's legacy system. For example, a "wrapper" could translate the legacy system's position data into a format that can be sent to and understood by the central risk management function. A financial firm with many legacy systems may face a higher cost of implementing a firmwide risk management function, because of the need for specialized "wrappers" for each legacy system, but the task is not impossible.

¹⁶ The magnitude of the approximation error involved in Value at Risk and stress test calculations based on mapping actual portfolio instruments onto a limited set of market risk factors could be large. For example, the error of mapping a 10-year bond with semiannual coupons onto a set of zero coupon bonds (say, 1, 2, 3, 5, 7, 9 and 10-year zeros) will vary with the shapes of the term structures of interest rates and volatility and the accuracy of the interpolations that must be made. As far as I am aware, this error has not been quantified in the risk measurement literature.

Many firms choose to replace their legacy systems rather than work around their limitations. A need to manage risks on a firmwide basis leads firms to insist that trading units have software that can interface smoothly with a central risk management function. Since many legacy systems lack this ability, many firms choose to replace them.¹⁷ As discussed briefly below, the European Union's Capital Adequacy Directive's requirements for firmwide measurement of market and credit risk have also led many firms to replace legacy systems.

3.4 "Smart" data structures

The use of "smart" data structures can significantly improve the cost-benefit tradeoff in favor of a rigorous (i.e., full revaluation Monte Carlo) risk measurement system. A "smart" data structure for storing financial transactions would have some or all of the following characteristics:

1. the financial instrument knows to which market risk factors its value is sensitive;
2. the financial instrument knows both a "more exact" and a "less exact" valuation method for itself;
3. the financial instrument knows what error is introduced at different times by its different valuation methods.

Exploiting such a "smart" data structure could significantly reduce the computational burden of revaluing each instrument in a portfolio N times to calculate a Monte Carlo Value at Risk.

If a financial instrument knows to which market risk factors its value is sensitive, it can avoid recalculating its value for some of the N Monte Carlo draws. If a DEM/USD currency swap is asked to revalue itself for several draws in which all market risk factors that would affect its value (presumably USD and DEM interest rates and the DEM/USD exchange rate) are identical, the value will be the same for all such draws and the computational burden can be reduced accordingly.

Many derivatives are typically valued using numerical methods. A simple example would be an American option, which could be priced exactly using a lattice as in Cox, Ross, and Rubenstein (1979). For risk measurement purposes, if there are many American options in the portfolio, it may be too time-consuming to value each American option on a lattice. Barone-Adesi and Whaley (1987) give an approximate analytic valuation method for an American option. A "smart" American option would know both valuation methods and would know how much error is introduced by the approximation at different times. This last feature would allow the risk management application to track how much uncertainty has been added to the estimate of firmwide Value at Risk by using "less exact" valuation techniques, as well as to set up a threshold for approximation error that would force the use of a "more exact" technique if a "less exact" technique gave a particularly bad

¹⁷ The information systems demands of risk management are driving thirty percent of UK banks to replace legacy systems, according to a recent survey. "Risk Management is Driving Banks to Replace Legacy Systems," *Risk Management Operations*, 16th December, 1996.

approximation for a certain set of changes in market risk factors.¹⁸ Our impression is that firms have begun to adopt "smart" techniques on an ad hoc basis, but their use has not yet become standard practice.

3.5 Credit risk measurement

There is no single measurement concept for credit risk, unlike Value at Risk for market risk, which has become widely accepted in the market. The current view of "best practice" for credit risk, as expressed in the Group of Thirty's (1993) report, is to measure both current exposure and potential future exposure, the latter calculated using statistical analysis of the future credit exposure and a broad confidence interval (two standard deviations). There are many ways to measure credit risk; current market practice can be divided into two groups: transaction methods and portfolio methods.¹⁹

Transaction methods compute a potential exposure for each transaction as notional principal times a multiplier that reflects the transaction's form (option, swap, etc.), maturity, and the inherent riskiness of the transaction's underlying market risk factor(s). The multipliers are calculated in advance for all combinations of transaction form, maturity, and riskiness of underlying to which the firm anticipates having exposures. The information systems requirements of transactions methods are small; to measure potential exposure, each transaction's multiplier must be looked up in the tables of multipliers that have been calculated in advance. Total credit risk with a counterparty is simply the sum of current and potential exposure of each transaction with the counterparty in the firm's portfolio.

Portfolio methods compute the potential exposure of all the firm's transactions with a counterparty at once, considering correlations between potential exposures of multiple transactions with the counterparty as well as netting arrangements. Portfolio methods of measuring potential credit risk exposure are conceptually close to methods of calculating Value at Risk to measure market risk, with the additional complication of identifying each transaction's counterparty and netting status. The information systems requirements would be similar to those for calculating Value at Risk. In particular, a similar choice of methodology (delta-normal or full revaluation Monte Carlo) exists for credit risk measurement, and the choice will have a strong influence on information systems requirements just as it does for market risk.

Portfolio methods are superior to transactions methods because they incorporate the correlation of credit risk exposure among all transactions with a counterparty and netting arrangement that serve to reduce credit risk exposure. Firms pursuing "industry best practice" either have adopted or hope to adopt a portfolio approach to measuring credit risk. Since the information systems

¹⁸ For example, if an at-the-money American option close to expiration were subject to unacceptably high approximation error, the numerical method would be used.

¹⁹ This discussion is based on two articles describing Citibank's credit risk measurement system: Picoult (1996) and Lawrence (1995).

requirements of such an approach are similar to the requirement for calculating Value at Risk, firms may enjoy a synergy if they can design a single information system to meet both needs.

4. Information systems and measuring aggregate market risk

To measure aggregate market risk, a central bank would be faced with many of the same problems vis-a-vis individual firms that firms are faced with vis-a-vis trading units when setting up a central firmwide risk management function. Since a central bank will presumably not be interested in a "centralized" solution, we must look to the decentralized solutions already discussed to explore this analogy further.

Compare the "decentralized" solutions for computing delta-normal and full revaluation Monte Carlo Value at Risk. In the former, each trading unit must supply information on its "deltas" with respect to a uniform set of market risk factors. To use this methodology to measure aggregate market risk, a central bank would have to specify a common set of market risk factors for all firms to use. Specifying a common set of market risk factors presents several problems. For those firms that use a delta-normal methodology to calculate Value at Risk, their set of market risk factors is likely to be different from the common set, requiring reprogramming of the mapping procedure. Those firms using a different methodology, one that does not require mapping, would have to devote information systems resources to devising a mapping scheme that is irrelevant for their own business needs.

In full revaluation Monte Carlo Value at Risk, each trading unit computes the change in portfolio value for each Monte Carlo draw. To use this methodology to calculate aggregate market risk, a central bank would have to specify a set of Monte Carlo draws for all firms to use. Firms using a full revaluation Monte Carlo methodology, whether centralized or decentralized, could compute their change in portfolio value for each Monte Carlo draw in the same fashion as they compute it when measuring Value at Risk with their own Monte Carlo draws. Firms using another methodology, such as delta-normal, will also be able to calculate their portfolio's change in value if they are able to do stress tests or scenario analyses. The common set of Monte Carlo draws could be processed as "stress scenarios" by such firms. A potential problem with this approach to measuring aggregate market risk is the burden of computing a large number of scenarios, on top of the calculations need to compute a firm's own risk measures.

4.1 Heterogeneous mappings and aggregate market risk

In section 3.2 above, we discussed the problem a firm faces in choosing a set of market risk factors to summarize its position data for computing Value at Risk and conducting stress tests and sensitivity analyses. Additional problems arise when combining the output of individual firms' risk management information systems to measure aggregate market risk.

If two firms use different mappings for the same position, they implicitly have different definitions of market risk and basis risk, making aggregation of market risk across firms problematic. For example, suppose two firms each have made a contract for delivery of NLG for USD at a certain exchange rate in 30 days. One firm has the appropriate NLG/USD risk factor in its set of market risk factors, so no mapping is required. The other firm, with little NLG/USD exposure, chooses to map its exposure onto a DEM/USD risk factor, implicitly categorizing changes in the correlation between NLG/USD and DEM/USD exchange rates as basis risk. Suppose a range of scenarios of exchange rate changes are provided to both firms, and a portfolio revaluation is requested for each scenario. Given the limitations of its risk management information system, the second firm will use only the DEM/USD exchange rate to evaluate the NLG/USD forward contract. The two firms will likely give different answers to the question "How much does the value of the forward contract change under each scenario?" because their mappings are different. The variance in the answers will be greatest for those scenarios that incorporate moves in the NLG/USD and DEM/USD exchange rates that are farthest from their historical correlation.

Similar problems would arise when using individual firm "deltas" with respect to particular market risk factors as a basis for measuring aggregate market risk. A firm's sensitivity to DEM/USD risk may incorporate other exchange rate risks that have been mapped onto the DEM/USD risk factor; if each firm has a different mapping, each firm's sensitivity will be measuring a different concept, again making aggregation problematic.

Different firms' business needs may lead them to choose different sets of market risk factors to build into their risk management information systems. This variation across firms creates a problem for aggregating measures of market risk across firms, since different sets of market risk factors imply different definitions of what is encompassed by "market risk." Investigation of how serious a problem, in practice, this variation across firms would create for measuring aggregate market risk and how to get around the problem is left as an exercise for further research.

4.2 Data on counterparty type

Depending on the use to which the data is to be put, data on aggregate market risk may need to be computed by counterparty type. An information system designed solely to measure market risk will not necessarily be able to produce a breakdown by counterparty of either cashflow mappings or changes in portfolio value from a Monte Carlo draw. However, an information system designed to jointly measure market and credit risk could track counterparty type. The merits of such systems have been discussed in industry trade journals, but it is unclear how many firms are implementing such systems.²⁰

²⁰ See "Risk Where Credit's Due", *Risk* 9:6 (June 1996) and "Together They Stand", *Risk Firmwide Risk Management Supplement* (July 1996).

5. Conclusion

To measure risk, a financial firm needs sophisticated information systems. The information systems must combine data from disparate trading units in a structured way to estimate the aggregate risk of the firm. In this paper we have outlined some of the issues firms face when setting up such systems. We have described two of the many risk management methodologies currently in use in the market, and shown how methodology and information system design interact.

Efforts by firms to construct information systems that measure their risk on a firmwide basis have enable us to consider the possibility of aggregating risk data across firms in a meaningful, timely way. If the many hurdles could be overcome, this could represent a revolutionary new way to construct market oversight information. To measure aggregate market risk, risk data must be combined across firms. Any such effort would face many of the same issues of information systems capabilities faced by firmwide risk management in an individual firm. While the issues are similar, the limited ability to impose coordination across firms raises new problems. How to overcome these problems is left as a task for future research.

References

- Barone-Adesi, G. and R.E. Whaley (1987): "Efficient analytic approximation of American option values", *Journal of Finance* 42 (June), p. 301-320.
- Bernardo, Antonio E. and Bradford Cornell (1997): "The valuation of complex derivatives by major investment firms: Empirical evidence", *Journal of Finance*, June.
- Carverhill, Andrew and Les Clewlow (1994): "Quicker on the curves", *Risk* (May).
- Cox, J.C., S.A. Ross and M. Rubenstein (1979): "Option pricing: a simplified approach", *Journal of Financial Economics*, 7 (October), p. 229-263.
- Group of Thirty (1993): "Derivatives: Practices and principles", Washington, DC, Group of Thirty.
- Lawrence, David (1995): "Aggregating credit exposures: The simulation approach", In *Derivative Credit Risk*, London, Risk Publications.
- Picoult, Evan (1996): "Measuring pre-settlement credit risk on a portfolio basis", In *Risk Measurement and Systemic Risk. Proceedings of a Joint Central Bank Research Conference*, Washington, DC, Board of Governors of the Federal Reserve System.
- Pierides, Yiannos (1996): "Legal disputes about complex interest rate derivatives", *Journal of Portfolio Management* (Summer), p. 114-118.
- Press, William H. et al (1992): "Numerical recipes in C: The art of scientific computing", Second edition, Cambridge University Press.
- Pritsker, Matthew (1997): "Evaluating value at risk methodologies: accuracy versus computational time", *Journal of Financial Services Research*, forthcoming.