

BANK FOR INTERNATIONAL SETTLEMENTS

**THE MEASUREMENT
OF AGGREGATE MARKET RISK**

A joint exploration by a group of central bank researchers

Basle
November 1997

The present publication is also available on the BIS World Wide Web site (<http://www.bis.org>).

© *Bank for International Settlements 1997. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISBN 92-9131-060-3

Table of contents

Preface	
Introduction and summary	
PART I - OVERVIEW	1
I. Goals and objectives	1
II. The group's research agenda	2
III. Dimension reduction and scenario specification	6
IV. The trade-off between accuracy and computational burden	12
V. Feedback, market liquidity and the role of information	13
VI. Concluding remarks	18
PART II - INDIVIDUAL RESEARCH PAPERS	
Generating market risk scenarios using principal components analysis: methodological and practical considerations - Mico Loretan, Federal Reserve Board	23
Term structure and volatility shocks - Anthony P. Rodrigues, Federal Reserve Bank of New York	61
Approximation of changes in option values and hedge ratios: how large are the errors - Arturo Estrella and John Kambhu, Federal Reserve Bank of New York	103
Information systems for risk management - Michael Gibson, Federal Reserve Board	115
Residual risk factors, portfolio composition and risk measurement - John Kambhu and Anthony P. Rodrigues, Federal Reserve Bank of New York	133
Liquidity risk and positive feedback - Matt Pritsker, Federal Reserve Board	145
Dynamic macro stress exercise including feedback effect - Tokiko Shimizu, Bank of Japan	163
Measurement of liquidity risk in the context of market risk calculation - Jun Muranaga and Makoto Ohsawa, Bank of Japan	193
Information collection and disclosure - M. Matthew Adachi, Bank of Japan and Patricia Jackson, Bank of England	215
The size of hedge adjustments of derivatives dealers' US dollar interest rate options - John Kambhu, Federal Reserve Bank of New York	223

PREFACE

This volume contains papers produced for the Euro-currency Standing Committee in a joint effort by researchers at several central banks. The papers address measurement of market risk, market dynamics, market liquidity, and the role that information plays in determining market outcomes in unsettled circumstances. The Committee believes that the research undertaken will be of interest to a wider audience, including market participants and the academic community. In publishing the papers, the Committee hopes to stimulate further research in these areas. The papers represent the views of the authors and not necessarily those of the central banks with which they are affiliated nor that of the Euro-currency Standing Committee.

In July 1996, the BIS published a report of a working group of the Committee, chaired by Shinichi Yoshikuni of the Bank of Japan, which recommended the establishment of a reporting system on activities in global derivatives markets. That reporting system is to be implemented in 1998. The Report recognised that data on derivatives positions, while indispensable for tracking changes in the size and structure of derivatives markets over time, would shed limited light on how overall portfolio values and market conditions might change in the face of price shocks. The behaviour of markets in the face of shocks has long been an area of fundamental central bank interest and responsibility.

When the papers in this volume were discussed by the Committee in May 1997, the Committee accepted the researchers' conclusion that this research did not establish an adequate technical basis or adequate justification for collecting aggregate market risk data. However, the Committee decided to encourage continuing work on other aspects of market behaviour addressed in these papers. In particular, in line with its mandate to monitor sources of potential instability in financial markets, the Committee will continue to encourage and review research on market functioning and price dynamics under stress.

Toshihiko Fukui,
Chairman, Euro-currency Standing Committee
Senior Deputy Governor, Bank of Japan

Introduction and Summary

The papers in this volume set out the results of a co-ordinated research effort by the Bank of England, the Bank of Japan and the Federal Reserve into issues to do with the measurement of aggregate market risk, a topic identified for further investigation in "Proposal for Improving Global Derivatives Market Statistics" (the Yoshikuni Report), issued by a working group of the Euro-currency Standing Committee in July 1996. The research considered whether data could or should be collected from market intermediaries on exposures to stress scenarios and then aggregated across whole markets.

The research covered the following questions:

- (a) the development of a methodology for designing stress tests which could be used in the collection of the aggregate market risk data;
- (b) wider questions regarding dynamic effects in periods of market stress, in particular feedback and market liquidity;
- (c) the benefits of collecting and disseminating aggregate market risk data.

The question of the design of a stress test methodology was important because if likely profits and losses under extreme scenarios were to be aggregated across firms, a methodology needed to be developed to generate those profits and losses. The Value at Risk methodology, which many of the large firms are developing, will deliver an exposure (given set parameters) but not whether that would be a loss or a profit.

A number of the papers look at whether principal components analysis (PCA) could be used to generate stress scenarios which could be applied to all the books held by intermediaries in a market. The conclusion reached was that it might be possible to use PCA to develop a stress test methodology, but that there were a number of issues regarding the dimension reduction - whether it would mean that risks inherent in the books were not captured.

The research took account of the likelihood that holders of financial assets could respond to a change in market prices by changing their desired portfolio composition, perhaps because they pursue dynamic hedging strategies. A framework to incorporate feedback trading and liquidity effects into an aggregate risk measurement process was considered. The framework relies on an as-yet-unproved ability to characterise the types of traders in a given market, in order to model both their feedback trading in response to particular scenarios of shocks and their impact on market liquidity.

The benefit of aggregate market risk information would depend on its ability to improve market functioning in times of market stress; that ability would depend on the information's accuracy, timeliness, and frequency, as well as upon whether mechanical feedback trading is large or small relative to investors' repositioning during market breaks. It was noted that distortions related to

signalling and moral hazard problems could arise if central banks play a role in formulating scenarios and in disseminating the results.

The group's overall assessment of the research can be summarised as follows. First, the research has not established an adequate technical basis or adequate justification for collecting data on aggregate market risk. It is not clear whether the data generated from the books of core intermediaries would convey much information about likely market dynamics in the aftermath of a shock, given the large fraction of wealth held - and trading accounted for - by institutional investors. While the group believes that the burden of producing useful aggregate market information will inevitably fall as new information system technologies are deployed, the cost of producing such information would still have to be evaluated relative to the benefits the information would provide. Second, there was agreement that the research effort had highlighted the importance of the roles played by dealers and non-dealers in market functioning, and that further research into market dynamics, market liquidity and market functioning in times of stress would be useful.

The volume is divided into two parts. Part I of the volume provides an overview of the research strategy that was pursued with respect to work on developing measures of aggregate market risk. This overview places the individual papers produced for the project within a conceptual framework. Part II of the volume contains the individual papers produced for the project.

PART I: OVERVIEW

I. Goals and objectives

A desire to better understand the dynamics of market liquidity motivated the research project. If a growing number of market participants are relying on market liquidity in their investment and risk management choices, then the robustness of that market liquidity would appear to be an important issue. One set of market participants who rely on market liquidity are those firms engaged in dynamic trading strategies, such as dynamic hedging or portfolio insurance. Previous research has highlighted the possibility that such strategies could, at times, have adverse repercussions for market functioning. The group set out to address the question whether the collection and dissemination of information about the potential scale of trading engendered by such dynamic risk management strategies would be feasible and useful.¹

One set of previous research efforts which provided valuable inputs into the group's efforts were the models that were developed to explicate the October 1987 stock market crash.² The important features of such models are the explicit treatment of hedging operations by market participants who are portfolio insurers, the explicit treatment of market liquidity, and the key role of information. For instance, if other market participants mistakenly believe that hedging-induced selling is driven by negative information, negative market price shocks can be amplified. This would not be the case if other market participants were aware of the magnitude of hedging demands generated by portfolio insurers. The crash models are useful devices to frame our discussions about the value of information on market risk profiles, since in those models such risk profiles provide information on potential hedging demands.

The research took into account the group's understanding of how firms manage their risk-taking activities. Based on conversations with market participants and bank supervisors, it was found that dealers' risk management activities continue to be, in many cases, organised by product desk. On the other hand, cutting-edge firms now measure their risk exposures on an integrated firm-wide basis. Such risk measures would incorporate, for example, the exposure to US dollar interest rates

¹ The work program was framed by a group of researchers from the Bank of England (Patricia Jackson), the Bank of Japan (Matthew Adachi, Jun Muranaga, Makoto Ohsawa and Tokiko Shimizu), the Federal Reserve Bank of New York (John Kambhu and Anthony Rodrigues), and the Federal Reserve Board (Allen Frankel, Michael Gibson, Mico Loretan and Matthew Pritsker). The group first met at the Federal Reserve Bank of New York in July 1996 to develop a research agenda that led to the papers in this volume. In December, the group met at the Federal Reserve Board to review the drafts of the research papers. In February, the secretary of the Standing Committee hosted a meeting of the group at the BIS with experts from the other ECSC member central banks. At this meeting, members of the enlarged group discussed the results of the papers, as well as the potential payoff from the pursuit of the broader research strategy.

² Two papers that we relied on most were Gerard Genotte and Hayne Leland, Market liquidity, hedges and crashes, *American Economic Review* 80 (1990), 999-1021 and Sanford Grossman, An analysis of the implications for stock and futures: Price volatility of program trading and dynamic hedging strategies, *Journal of Business* 61 (1988), 275-298.

originating from all cash and derivatives market transactions, which could include non-trading exposure such as loans and deposits. Over the medium term, we expect that it will become standard for dealer firms to employ firm-wide risk measurement and risk management techniques. An attempt was made to reflect these expectations in the design of the research strategy, by not confining the analysis of aggregate risk measurement techniques to what would be possible with the information systems in place today. In framing its work program, the group also appreciated the need to evaluate the trade-off between accuracy and computational burden.

The remainder of this overview is organised as follows. Section II provides a summary of the research strategy that was pursued with respect to work on developing measures of aggregate market risk. The individual papers that were produced for the project are placed within the conceptual framework. Sections III, IV, and V discuss, in more detail, the specific areas of research with which the papers are concerned. Section VI provides the group's assessment of the progress achieved.

II. The group's research agenda

The group's research effort began with the recognition that some measures of risk, while useful for individual firms, cannot be aggregated. For example, Value at Risk, the leading market risk measurement methodology, summarises the size of an exposure but not the its direction, as the example in Chart 1 overleaf illustrates. Two hypothetical firms hold offsetting positions of \$1 million each in the S&P 500 equity index, as shown in Panel A of the chart. Panel B shows two firm-level market risk measures. Each firm's Value at Risk, using the common set of assumptions given at the bottom of the chart, is equal to \$15,100. With a scenario-based market risk measure, each firm summarises its risk by computing the change in the value of its portfolio for a common set of five scenarios. The five scenarios are a change in the index from -2 percent to $+2$ percent, in increments of 1 percent. The resulting risk profiles summarise both the size and the direction of each firm's exposure.

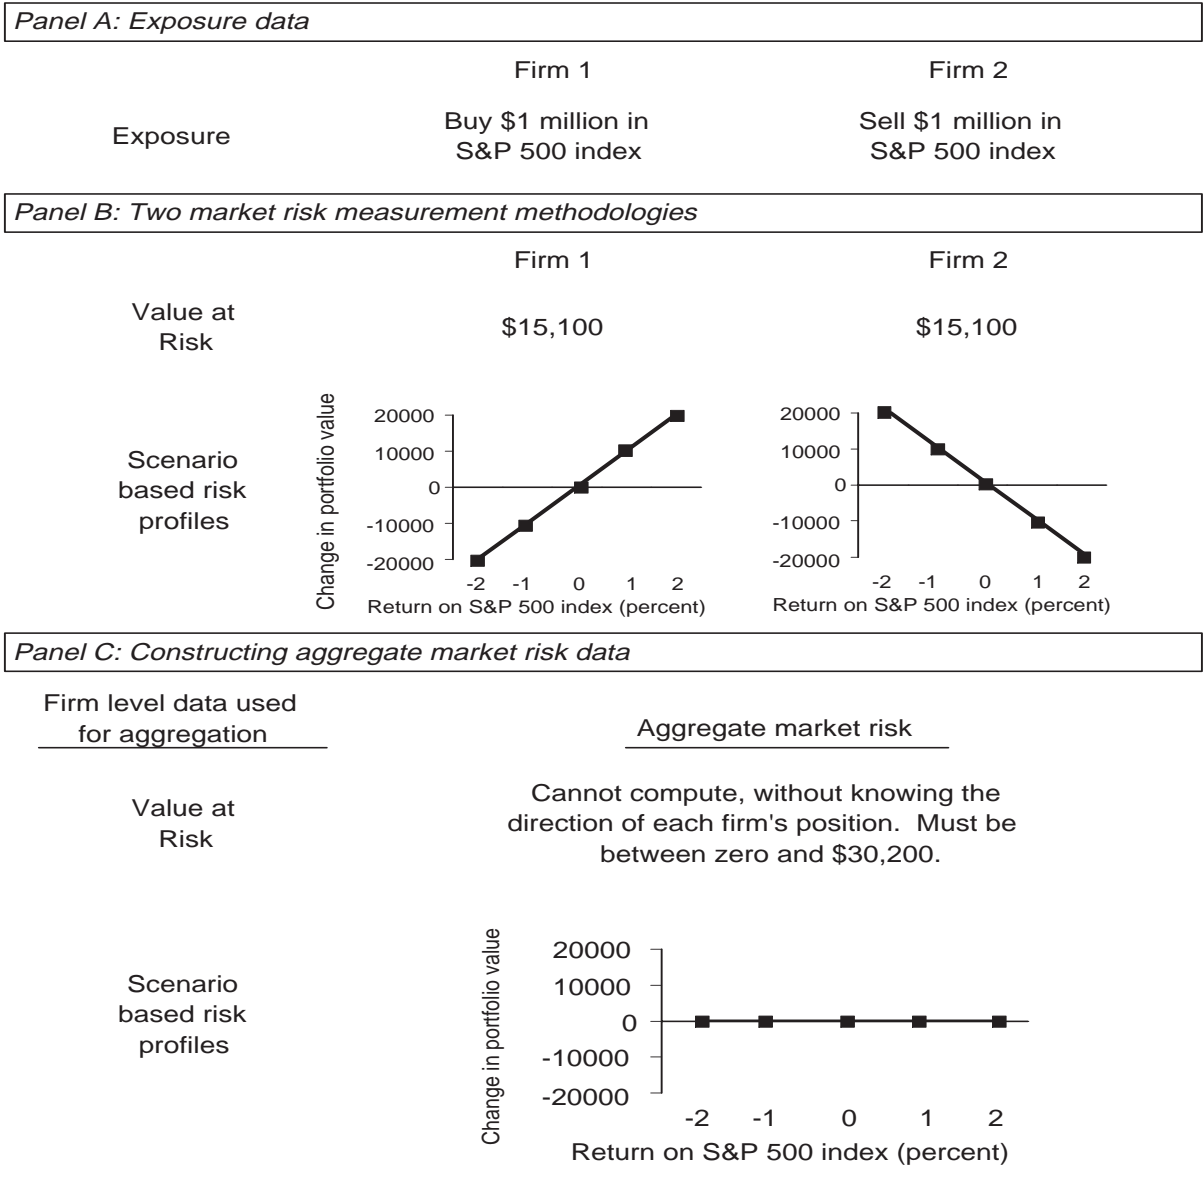
Panel C of Chart 1 shows how individual firm market risk data can be used to measure aggregate market risk. Because an individual firm's Value at Risk gives no indication of the direction of its exposure, the sum of the individual firms' Values at Risk can only provide an upper bound on aggregate Value at Risk. In the example in the chart, the upper bound is particularly uninformative since the firms' exposures exactly offset one another and aggregate market risk is zero. In contrast, because a scenario-based method conveys both the size and direction of exposure, an aggregate market risk profile can be easily computed as the simple sum of the individual firms' risk profiles. As shown in the graph in Panel C, the sum of the individual firms' risk profiles is zero for all five scenarios.

Chart 1 also can be used to illustrate the limitations of an aggregate measure of market risk as a stand-alone indicator of potential market stress. Assume that firm 1 is strongly capitalised and

firm 2 is critically undercapitalised. This assumption has the straightforward implication, all other things being equal, that the potential for systemic stress would be greater for increases in the S&P 500 index than for declines. The ranking of possible outcomes for systemic stress is not conveyed by the computed aggregate measure of market risk, which equals zero for all values of S&P 500 index returns.

Chart 1

Comparison of Value at Risk and scenario-based measurement of market risk

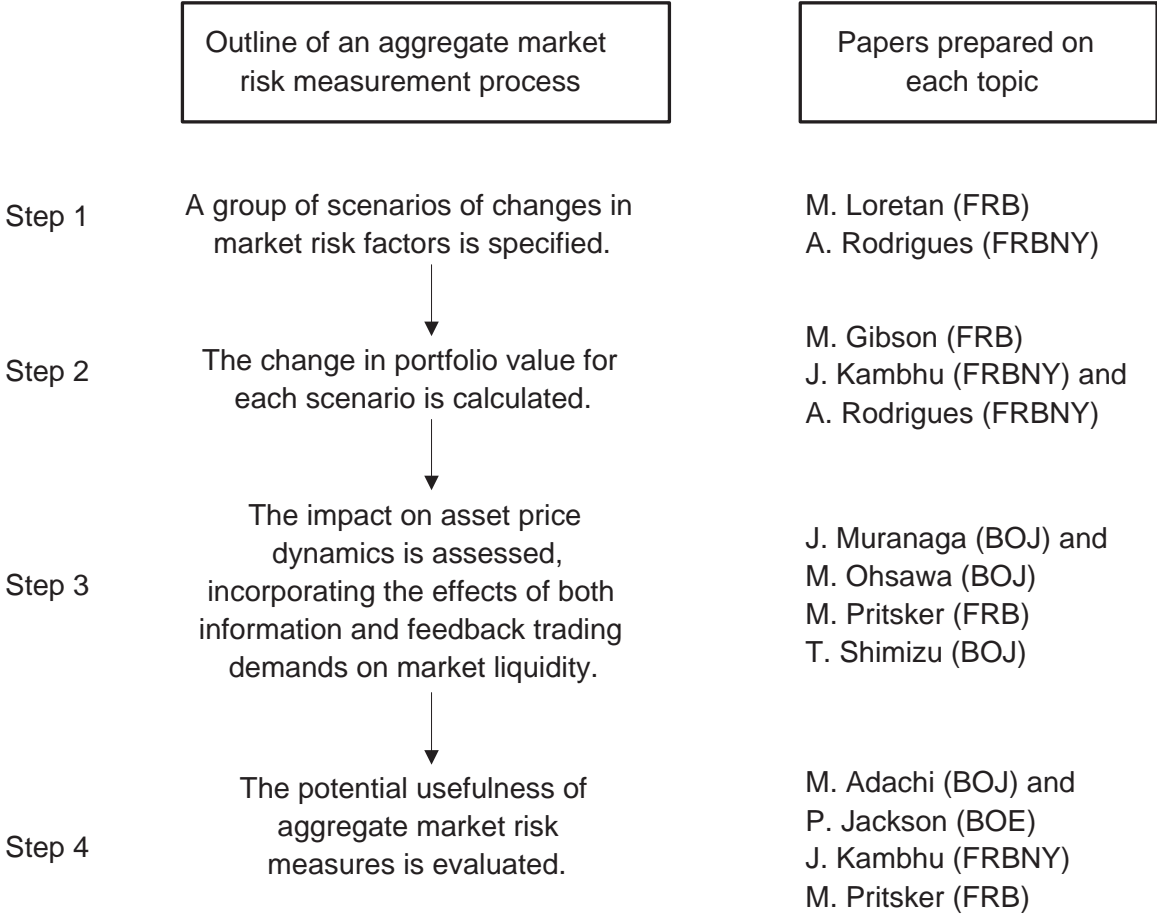


Assumptions:
 For Value at Risk: 1. Returns on the S&P 500 index are normally distributed. 2. Standard deviation of daily returns is 0.65 percent. 3. Time horizon is one day. 4. Confidence level is 99 percent.
 For scenario-based method: the five assumed scenarios are changes in the S&P 500 index of -2 percent, -1 percent, zero, +1 percent, +2 percent.

Against this background, the group's research strategy focused on a scenario-based measure of market risk, as outlined in Chart 2 below. To aggregate a scenario-based risk measure across firms, each firm must evaluate its risk with a common set of scenarios and comparable valuation methodologies. Thus, the group's research efforts began with investigations of scenario specification and valuation techniques. Firms that use a scenario-based risk measurement methodology face many of the same issues that the group studied in steps 1 and 2.³ Steps 3 and 4 reflect the group's concern with overall market functioning. Here the group's research goes beyond what individual firms currently do to measure their own market risk. Step 4 considers whether the availability of aggregate market risk data would matter for market outcomes and whether the benefits of collecting the data would outweigh the costs.

Chart 2

Overview of the group's research agenda




³ For a description of the scenario-based approach to risk measurement in use at Sakura Global Capital, see Farshid Jamshidian and Yu Zhu, Scenario simulation model: theory and methodology, *Finance and Stochastics* 1 (1997). A less technical description can be found in Farshid Jamshidian and Yu Zhu, Scenario simulation model for risk management, *Capital Market Strategies* 12 (December 1996). For another scenario-based approach, see Jon Frye, Principals of risk: Finding Value-at-Risk through factor-based interest rate scenarios, unpublished working paper, June 1996, Nationsbanc-CRT.

Step 1 in the process outlined in Chart 2 is to specify a group of scenarios of changes in market risk factors. The papers by Loretan and Rodrigues explore a methodology for specifying scenarios using statistics to summarise the movements and co-movements of financial asset prices with a small number of unobservable factors. Scenarios could be based on movements of a specified probability (e.g., 90, 95, or 99 percent confidence level) in these underlying factors. These papers reflect an interest in a scenario design mechanism that could measure risk profiles, based on statistical reasoning. In contrast, firms that use ad hoc scenarios of changes in market risk factors for internal risk management purposes typically do so to get a narrow view on their exposure to particular, extreme market situations as a supplement to a Value at Risk measure, rather than to measure their comprehensive risk profile.

In Step 2, each firm evaluates its own portfolio's value for each scenario. The primary issue to be addressed at this step is the cost-benefit trade-off between the accuracy and the computational burden of potential risk measurement methodologies. Chart 3 below represents feasible combinations of accuracy and computational burden as points on a curve. Greater accuracy (moving up on the chart) can be achieved by incurring a greater computational burden (moving to the right). For example, in a scenario-based aggregate risk measurement methodology, increasing the number of scenarios sampled would improve the accuracy of risk estimates, representing a movement from A to B on Chart 3.

Chart 3

The trade-off between accuracy and computational burden



Several of the group's papers address the trade-off between accuracy and computational burden. The paper by Estrella and Kambhu analyses how the accuracy of the valuation of options portfolios would improve when changing from an approximate revaluation methodology to a full revaluation methodology. Such a change would represent a move up and to the right on the curve in Chart 3. The paper by Gibson studies how the design of firms' internal risk management information

systems is influenced by the trade-off, as well as the broader question of what the current (and future) capabilities of firms' internal risk management information systems have to teach us about aggregate risk measurement. The paper by Kambhu and Rodrigues discusses an issue related to portfolio valuation that arose in the course of the group's research efforts: How does the choice of a methodology for scenario construction affect the accuracy of a firm's valuation of its portfolio?

Step 3 recognises that holders of financial assets could respond to a change in market prices by changing their desired portfolio composition, perhaps because they pursue dynamic hedging strategies. Depending on market liquidity conditions, these "feedback trading" effects could amplify or dampen a stressful shock to financial markets. The paper by Pritsker outlines the theory of feedback trading, market liquidity and their interaction. The paper by Shimizu explores how feedback trading could be incorporated into an aggregate market risk exercise and provides an illustrative example of how the feedback trading might be modelled. The paper by Muranaga and Ohsawa proposes modified market risk measures reflecting intra-day liquidity patterns as well as the market impact of trading activity and shows to what extent the quantified liquidity effects could affect conventional Value at Risk measurements of market risk.


Having assessed the impact of information, feedback trading and market liquidity on asset price dynamics in Step 3, the benefit of information on aggregate market risk to market participants and central banks is assessed in Step 4. The description above of hedging behaviour in "crash" models showed one potential benefit of such information. The theoretical argument by Pritsker about "sunshine trading" underlines the role of information in market functioning. The paper by Adachi and Jackson formulates implications for market functioning of information disclosures and considers whether data on aggregate market risk would convey information on likely market stresses in a crisis. The theoretical research into the benefit of information is complemented by Kambhu's empirical research estimating the size of hedge adjustments stemming from dealers' US dollar interest rate options portfolios. His results enable us to gauge what the benefit of aggregate market risk data would have been in April 1995.

III. Dimension reduction and scenario specification

The computational burden of a scenario-based approach to market risk measurement increases with the number of scenarios needed to capture the risk. The dimension of the scenario construction problem can be reduced by identifying a small number of market risk factors that drive most of the variation in market prices and volatilities. Scenarios could then be specified in terms of shocks to the few driving market risk factors. For example, Chart 4 overleaf graphs two hypothetical market data series. The hypothetical data series are positively, but not perfectly, correlated. Dimension reduction could be accomplished by representing both market data series with a single risk factor, shown in the chart as a solid line. Scenarios of shocks could be specified as points on the line.

Note what is lost with dimension reduction: only data points actually on the line can be represented with the single risk factor. Other points, not on the line, can only be approximately represented.

Chart 4
**Two hypothetical market data series;
the solid line represents the most important underlying risk factor**



The dimension reduction technique used in the papers by Loretan and Rodrigues is principal component analysis (PCA). PCA has been widely applied by financial firms, for portfolio management as well as for risk management.⁴ PCA is a statistical technique to summarise the variation among several variables. PCA identifies explanatory "factors," linear combinations of the original data, that explain as much of the variance in the data as possible. PCA proceeds sequentially, with each subsequent factor after the first required to be uncorrelated with all previously identified factors. Each subsequent factor captures a smaller fraction of the variance of the data than the preceding factor.

The goal of the exploration of PCA was to investigate its use as a tool to summarise the variation in asset prices across countries. For example, researchers who study movements of the term structure of interest rates in a single country have used PCA to identify three factors that explain most of the variation in the term structure. These three factors typically look like a "shift," "tilt," and "twist" of the term structure.⁵ Loretan and Rodrigues extended this approach to look across countries as well.

⁴ For an application of PCA to portfolio management, see Robert Litterman and Jose Scheinkman, Common factors affecting bond returns, *Journal of Fixed Income* 1 (1991), 54-61. For applications to risk measurement, see Ronald Kahn, Fixed income risk modelling in the 1990s, *BARRA Newsletter*, Winter 1995; Jamshidian and Zhu, op. cit.; and Frye, op. cit.

⁵ See Litterman and Scheinkman, op. cit.

Loretan focuses on general statistical issues arising with PCA, using daily data on spot exchange rates, equity prices, and interest rates for several countries. Rodrigues focuses on term structures of interest rates, within and across countries. In both studies, the main conclusion is that PCA may indeed be used to condense the multivariate properties of certain aspects of these datasets into a small number of risk factors. For example, Loretan finds that in a set of eight spot exchange rates, the most important risk factor—identified as the first principal component of the data—alone accounts for 72 percent of the total variance of the normalised data. The second risk factor accounts for an additional 13 percent. These results suggest that a fourfold reduction in dimensionality—from eight exchange rates to two risk factors—can be achieved at the loss of 15 percent of the covariation among the eight exchange rates, with an attendant loss of accuracy in portfolio valuation. Similarly, Rodrigues reports that the first few principal components of a 63-dimensional vector of interest rates in seven countries explain a large majority of the joint variation of longer-maturity bond returns.

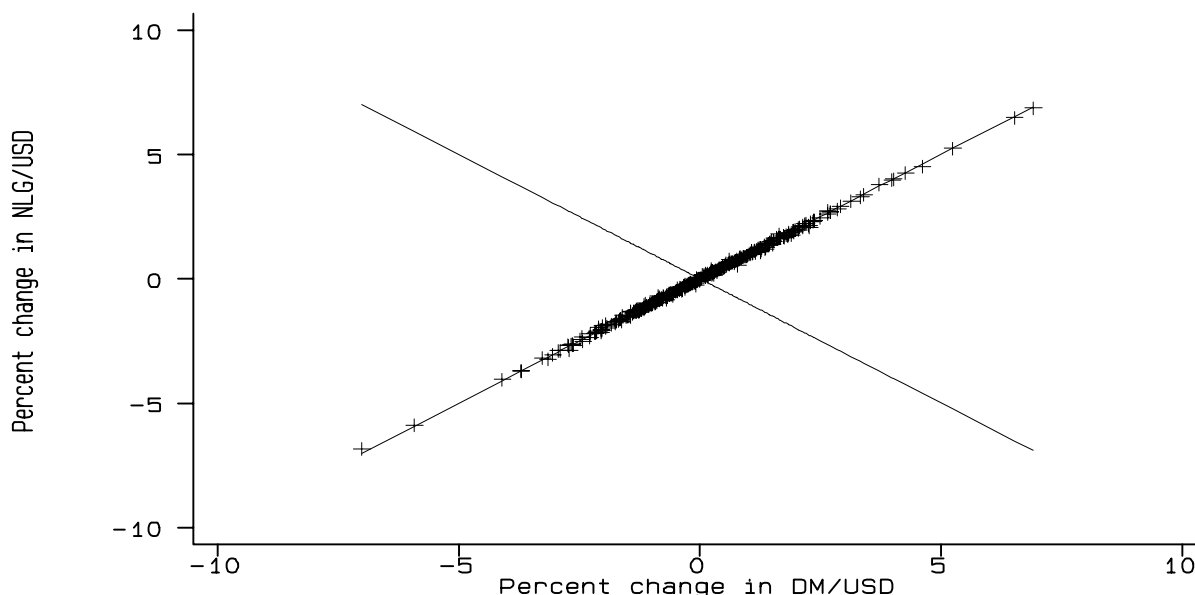
Both authors show that the amount of dimension reduction varies across instruments. For example, in Loretan's set of short-term interest rates in nine countries, the first two risk factors explain only 38 percent of the total variance, suggesting that not much dimension reduction would be possible for those series. Rodrigues shows that while a few principal components account for much of the variation in long-term bond returns across countries, short-term interest rates of different countries are poorly explained by a few common components. Whether the inability of PCA to summarise certain sets of data series with a few risk factors is an important shortcoming will be discussed below, when we consider the possible uses of PCA to measure a portfolio's market risk.

The motivation for examining PCA's ability to do dimension reduction was to reduce the computational burden of a scenario-based measure of market risk by reducing the number of instances for which it is necessary to perform portfolio valuations. Using the distribution of risk factors identified by PCA, it would be straightforward to develop scenarios that may be given probabilistic interpretations such as "10% (5%, 1%) tail events." Such scenarios would have a consistent meaning over time, and no subjective judgements would enter into the scenario construction process. As a result, second-guessing as to the motivation for scenario construction at any point in time could be reduced. While such second-guessing will always be an issue, it would be particularly important to address if, at some point, central banks were to be involved in putting together aggregate risk measures.

The group's discussion highlighted several issues and problems with PCA. The most significant question was whether the loss of accuracy in a risk measurement technique that uses dimension reduction would be so large and uncertain as to be unacceptable. The following example illustrates the problem. The US dollar/Deutsche Mark exchange rate and the US dollar/Dutch guilder exchange rate are highly correlated, as Chart 5 shows. A principal component analysis of the two series shows that a single risk factor can explain more than 99 percent of the total variance, suggesting that a 50 percent reduction in dimension (from two to one) could be achieved with little loss of explanatory power by summarising the movements in both series with a single risk factor (the positively sloped line in Chart 5).

Chart 5

**Weekly changes in the DM/USD and NLG/USD exchange rates,
January 1990 through February 1997, shown as +.
The two lines represent the two principal components.**

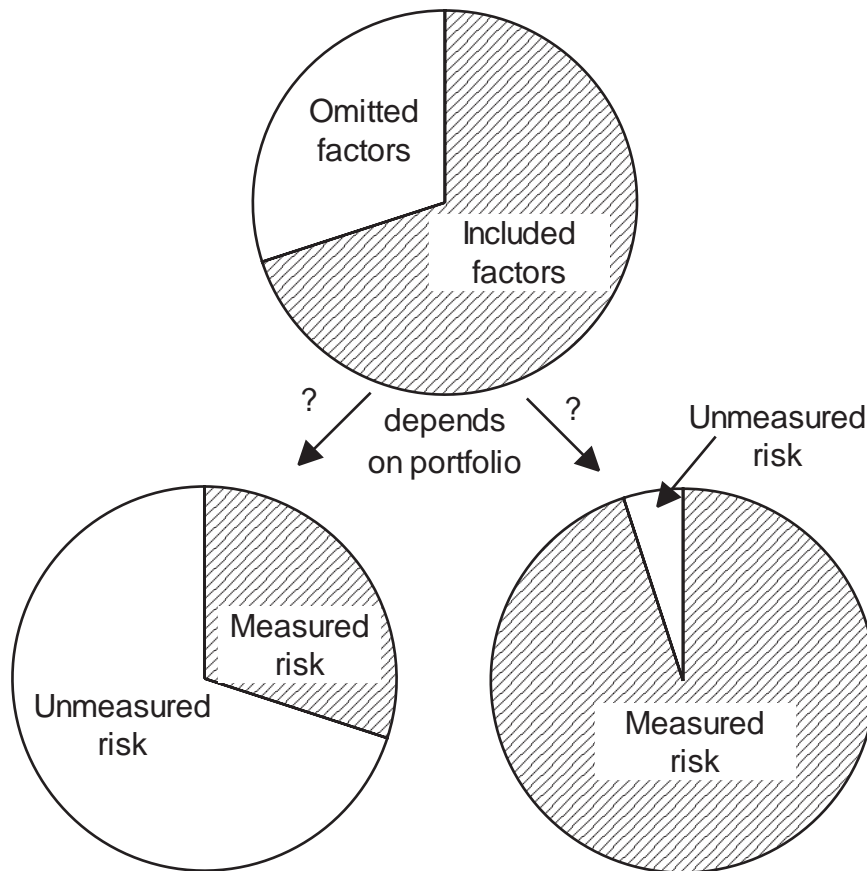


Continuing the example from Chart 5, scenarios could be constructed from movements in the single risk factor by choosing points along the positively sloped line. However, measuring market risk this way will not capture the risk that the two exchange rates could move in opposite directions. If a firm has exposure to the DM/guilder cross rate, dimension reduction will understate the firm's market risk. Research undertaken by Kambhu and Rodrigues in response to this problem suggests that the problem will be more severe and prevalent when taking a position in the omitted risk factors can increase a portfolio's expected returns. Even worse, if firms have a tendency to hedge the "obvious" risks (the first risk factor in the DM/guilder example) but take exposures to the less obvious risks (the second risk factor), dimension reduction could lead to a large, systematic understatement of the true market risk.

Chart 6 on the next page illustrates this point. The upper circle represents the entire variation of asset prices, of which only a fraction is captured due to dimension reduction. The two lower circles represent the entire portfolio risk, of which only a fraction is measured because of dimension reduction. Depending on whether the portfolio's exposure to the omitted risk factors is large or small, the amount of unmeasured risk due to dimension reduction can be large, as in the left-hand circle, or small, as in the right-hand circle. Chart 7 shows the effect of omitting fewer factors (i.e., doing less dimension reduction). (The cross-hatched area in each circle shows the change from Chart 6). Including more factors shrinks unmeasured risk, but the amount of unmeasured risk could still be large or small. This issue is another manifestation of the trade-off between accuracy and computational burden that was mentioned above and will be more thoroughly discussed below.

Chart 6

Dimension reduction and the accuracy of risk measurement



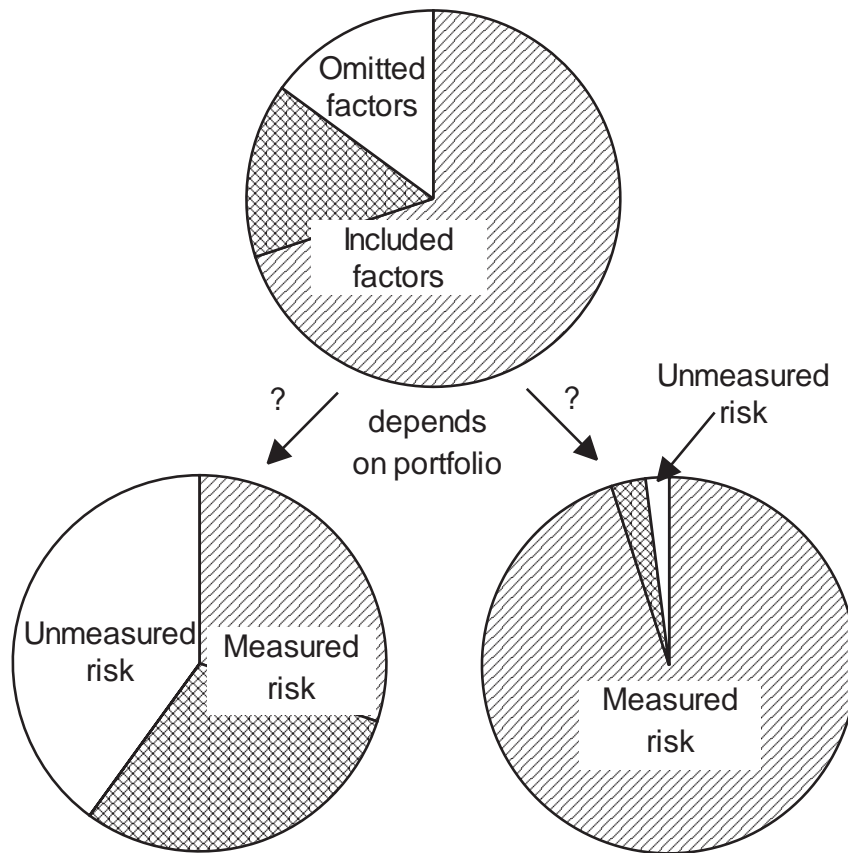
The degree of dimension reduction that should be done in any particular risk management scheme and the loss of accuracy of the resulting risk measures are open questions. Firms that use PCA for their own risk management purposes have faced this trade-off and have chosen to do some dimension reduction to reduce their computational burden in exchange for lower accuracy. However, firms face a relatively favourable trade-off because they know their portfolio, and they know in which directions they can afford to do dimension reduction with a limited loss of accuracy.

To illustrate this point further, assume that the two circles at the bottom of Chart 6 and 7 represent two firms with different portfolios. Because the firms have different portfolios, the same amount of dimension reduction leads to different amounts of measured and unmeasured risk at the two firms. The firm whose portfolio risk is represented by the circle on the left would be more likely to conclude that the increase in accuracy from doing less dimension reduction (i.e., moving from Chart 6 to Chart 7) is worthwhile.

The group had originally expressed an interest in a mechanical method of dimension reduction, retaining risk factors that account for a "large enough" fraction of the variance of asset

Chart 7

**Dimension reduction and the accuracy of risk measurement;
less dimension reduction leads to less unmeasured risk**



prices. As the above discussion shows, a mechanical procedure may drop risk factors that are important sources of variation in the value of some firms' portfolios. As the paper by Adachi and Jackson points out, the value of data on aggregate market risk depends strongly on its accuracy. If further research into measuring aggregate market risk is undertaken, this is an important issue that must be addressed: How can dimension reduction procedures be guided to achieve the highest possible accuracy for a given number of scenarios?

An alternative scenario construction technique, in widespread use at firms but not investigated by the group, is to use historical data to generate scenarios. With this technique, a firm would measure its market risk by revaluing its portfolio for each day's change in market prices over some historical time period. Assuming that the future distribution of market risk factors is like the past, the distribution of future changes in a portfolio's value can be approximated by the distribution of changes in portfolio value that would have occurred had the portfolio been held over the specified historical time period. This technique has the advantage of avoiding dimension reduction and its attendant problems, discussed above.

Risk measurement based on historical simulation has its own problems, which we highlight as potential topics for future research. A serious problem with aggregating historical simulation data

across firms is asynchronous data. Different firms' databases may record the same data series at different times during the day. For example, a Japanese bank may record the yen/dollar rate at the end of the day in Tokyo, while a US bank may record it at the end of the day in New York.⁶ As a consequence, the two firms' historical simulation results would not be directly comparable. Other issues to be investigated on the historical simulation approach include the increase in computational burden, with an unknown but potentially small gain in accuracy, and the difficulty in interpreting the results without the underlying factor structure that would be imposed on scenarios by PCA. Still, given the concerns about dimension reduction generated by the group's discussions, historical simulation is an alternative to PCA that could be investigated.⁷

IV. The trade-off between accuracy and computational burden

As noted, a key issue for any risk measurement scheme is the trade-off between the accuracy of the risk measurements that are produced and the burden of computing the risk measurements. Estimating the risk of non-linear positions with full revaluation methods, rather than approximate methods, can improve accuracy. Estrella and Kambhu's paper analyses the accuracy of a delta-gamma approximation⁸ to the change in value of interest rate options in response to a large change in interest rates. They find that a delta-gamma approximation produces large errors for options that are far out of the money and close to maturity, and small errors everywhere else.

Gibson's paper describes a firm's goals in setting up a risk management information system, how different risk measurement methodologies trade off accuracy and computational burden, the current "state-of-the-art" in risk management information systems, and likely future developments in risk management information systems. He identifies several ways firms can improve the accuracy of their risk measurements at the expense of a greater computational burden. For example, by measuring risk on a firmwide basis, rather than trading desk by trading desk, accuracy can be greatly improved by taking account of the correlation between the risks taken by different desks. To recast the example from Chart 1, if two desks have equal but opposite positions in a particular market, aggregate market risk is zero while each trading desk's market risk is greater than zero. Firmwide risk management comes at a cost: either each desk must pass its position data up to a centralised risk manager who calculates a measure of firmwide market risk, or each desk must produce additional risk data that can

⁶ JP Morgan's RiskMetrics™ methodology has studied the effects of asynchronous data on the estimates of correlations between two data series. To reduce this problem, JP Morgan has moved much of its data collection to 4:00 pm London time. See JP Morgan/Reuters, RiskMetrics™ Technical Document, Fourth Edition, 1996, p. 184-196. (available on the Internet at

<http://www.jpmorgan.com/RiskManagement/RiskMetrics/RiskMetrics.html>)

⁷ Other scenario construction techniques that have been suggested include using implied volatilities from option prices or time-varying volatility models.

⁸ An option's value depends on the price of its underlying risk factor in a non-linear fashion. A delta-gamma approximation uses a quadratic function to approximate this non-linear relationship.

be aggregated at the firmwide level. This cost is paid largely in terms of computer hardware, networking, database storage, and programming effort.

Often a firm's motivation for not implementing firmwide risk management is the difficulty of accessing the position data held in each trading unit's computer systems. The difficulty arises because trading units may not have had an incentive to design their software systems with an eye towards providing data for firmwide risk management. In the field of information systems, software systems whose age or design prevents them from doing what new, well-designed software can do are referred to as "legacy systems." Because the problem of legacy systems is widespread, solutions are widely available and the cost of dealing with legacy systems is well-understood.

Over time, technical progress will allow firms to achieve the same level of accuracy with less computational cost, or greater accuracy with no increase in computational burden. On Chart 3, technical progress shifts up the curve representing the feasible combinations of accuracy and computational burden, moving a firm from a point like A to a point like C. Consequently, as computing costs continue to fall, many firms could be expected to improve the accuracy of their risk measurements by increasingly using full revaluation methods and centralised risk measurement. Of course, superior risk measurement will not necessarily reduce risks, as the technology that reduces the risk in a given trading strategy could lead some to adopt riskier trading strategies.

In large part, the development of derivatives markets and their role in intermediating risks have been driven by advances in computer technology. Further advances have the potential to continue to influence the nature of financial intermediation and the risk management services that banks and securities firms offer. Future research on this impact on financial intermediation and trading activity may be useful to central banks in executing their responsibility for maintaining financial stability.

As advances in computing and information technology continue to enhance firms' risk management capabilities, such changes might also affect considerations regarding the type of information market participants and central banks would like to see. Information about market prices, market liquidity, transaction volume, or market size that market participants themselves desire in the future may go beyond currently available information.

V. Feedback, market liquidity and the role of information

Members of the group undertook research into feedback trading, market liquidity, and the role of information to improve our understanding of how stressful market conditions could impair market functioning. Data on firms' risk exposures to scenarios could be used, along with the scenarios themselves, as inputs to a model of feedback trading. The feedback trading demands of all market participants would combine with the distribution of information among market participants to form new equilibrium asset prices. Different information could lead to different price outcomes. Results in

this area would enable one to assess the benefits of collecting aggregate market risk data and providing it, in some form, to market participants.

The existence of trading strategies which lead to market feedback is well-known. To give an example of such a strategy, consider a dealer who has written a put option on a stock. The dealer could hedge by selling short a number of shares equal to the option's "delta." Since the option's delta varies with the price of the underlying stock, the amount of the hedge must be continually adjusted, making this a dynamic hedging strategy. Because the option's delta rises as the stock price falls, dynamic hedging leads to additional selling of the stock during price declines. If the selling is misconstrued as motivated by negative information on the firm's prospects, it could lead to additional downward pressure on the stock price. Besides dynamic hedging, other trading strategies that might lead to market feedback include stop-loss trading and limit orders.

While the existence of such strategies has been well understood, less is known about their empirical relevance or how they could be incorporated into an aggregate market risk measurement process. The group's research efforts were directed in this direction. Pritsker's paper outlines how feedback trading strategies could be incorporated into a model of price determination. Shimizu's paper presents a technique for summarising the strategies of trading agents using artificial intelligence techniques, recognising that those strategies are complex and heterogeneous.⁹ Since information on the actual behaviour of traders is not available, she implements her technique using simulated trading data derived from laboratory trading sessions. Further studies will be required to implement it with real data. Commercial applications of artificial intelligence to financial market trading may also have the potential to be adapted to these uses.

In the second half of her paper, she describes how a population of artificially intelligent trading agents could simulate feedback trading, which could then be incorporated into a model of price determination such as the one outlined by Pritsker. One would need to be able to classify trading agents into a limited number of categories in order to represent the real-world population of trading agents with a group of artificially intelligent agents. The process of categorising trading agents' strategies is analogous to the way mutual funds are categorised, by themselves or others (as "growth fund," "income fund," etc.) but is much less well developed. Since mutual fund categories stem from a desire to evaluate relative performance among funds with similar strategies, the process of categorising trading agents' strategies will benefit from efforts to benchmark performance, presumably motivated by firms' desire to rationalise the structure of compensating risk-taking on their behalf. It is still an open question whether these efforts will yield the information required to model empirically the interactions of trading strategies.


Market liquidity is another issue considered to be crucial to market functioning, especially in stressful situations. As Chart 8 cited from the paper by Muranaga and Ohsawa illustrates, liquidity

⁹ Artificial intelligence is the science of designing computer systems to mimic the outcomes of human decision-making.

risk can be interpreted as a sum of the variability of execution cost (cost of immediacy) and that of opportunity cost (cost of waiting). The execution costs, comprising bid-ask spreads and the price impact of trading activity - so-called "market impact" - decrease with the time needed to complete an intended transaction, while opportunity cost tends to increase with execution time. Both costs are likely to increase considerably in a stressful situation, which underlines the importance of incorporating liquidity risk into aggregate market risk measurement. The paper by Muranaga and Ohsawa proposes modified market risk measures reflecting intra-day liquidity patterns as well as market impact, and demonstrates, by applying those risk measures to the Japanese equity market, to what extent the quantified liquidity effects could affect conventional Value at Risk measurements of market risk. In reflecting on the implications of their work for future research on aggregate market risk issues, Muranaga and Ohsawa emphasise the need to address liquidity issues using a dynamic framework that allows for interaction between trading activity in response to a shock and subsequent price movements.

The theoretical literature on price formation and the role of information is also at a fairly primitive level, due to the complexities involved in modelling interactions among investors with different information and different motivations for trading. The group's research efforts in this direction made little progress over the current state of the literature, which can be summarised with the following example from a leading graduate textbook on the subject:¹⁰

Chart 8
Components of liquidity risk



¹⁰ The example is paraphrased from page 204 of Maureen O'Hara, *Market Microstructure Theory*, Blackwell Publishers, Inc.: Oxford, 1995.

Compare price determination in a market where the entire order book, including stop-loss orders,¹¹ is public information with a market where the order book is not public so traders cannot distinguish between ordinary sell orders and stop-loss-generated sell orders. Following negative news, informed traders sell and the price falls in either market, triggering some stop-loss orders. If the order book is public information, traders, knowing that stop-loss orders do not reflect any additional negative news, will not lower their price. If the order book is private, on the other hand, traders cannot tell whether the new orders are stop-loss orders or information-based; prices continue to fall as a result.

A real-world example of information's role in market functioning shows why further research in this area could be fruitful. The practice of "sunshine trading" refers to trades that are not motivated by any new information and which are publicly announced to the market. Sunshine trades should not move market prices, since they contain no new information about fundamental values. Holding the number of trades and their sizes constant, publicly announcing sunshine trades implies there are fewer trades that cause prices to change, so the variance of prices falls. Market liquidity could also be improved by sunshine trading, since a lower variance of prices may make it less costly for risk-averse agents to enter the market to speculate.


Another set of issues regarding the role of information in an aggregate market risk exercise concerns the disclosure of aggregate market risk information and the role of central banks. The paper by Adachi and Jackson surveys these issues. The authors note that the usefulness of aggregate market risk data would depend on its accuracy, as discussed above, as well as on its timeliness and frequency.

Adachi and Jackson also address possible distortions that could arise if central banks play a role in disseminating aggregate market risk data. If a scenario-based risk measurement methodology is used, and if the methodology underlying the scenarios is not clearly explained and free from judgement, market participants might interpret the scenarios as a signal from central banks considering their current concerns. Also, as with any central bank intervention in private markets, the potential for moral hazard exists. The authors discuss measures to minimise such distortions, which might include the central banks' commitment to the scenario generation process.

The main issue raised in the discussion of the research on feedback trading, market liquidity, and the role of information was whether an aggregate risk measurement process limited to core dealers would capture the relevant volume of dynamic trading behaviour. For example, if non-dealers conduct the bulk of the dynamic hedging, the scope of the aggregate market risk measurement process would have to be broadened to include non-dealers. Members of the group differed in their conjectures as to how broad the process would have to be. The breadth of the aggregate market risk measurement process can be thought of as a point along the following continuum (illustrated

¹¹ A stop-loss order is an order to sell when the market price falls below a certain level.

overleaf), stretching from regulated financial intermediaries through core dealers, through financial firms in general, ending with all financial market participants including end users such as pension funds and mutual funds:




Some members of the group felt that including core dealers would be sufficient to capture the relevant volume of dynamic trading strategies. These members cited the possibility that technological change in financial markets will lead to a greater concentration of price risk intermediation in core dealer firms. While the desire on the part of end users to acquire portfolio insurance may not change, economies of scale may make it more attractive to purchase that insurance from a core dealer rather than provide it in house with dynamic trading strategies. Other members of the group felt that a focus on core dealers would be too narrow, missing a large fraction of the positive feedback trading that goes on. All members agreed that existing empirical research gave little guidance on this question. However, data on counterparty breakdown of dealers' option positions, including as of mid-1998, semiannual statistics on global derivatives markets (Yoshikuni data) and from firms' public disclosures, may prove useful to assess the scale and distribution of price risk intermediation.

As a complement to the theoretical papers of Adachi and Jackson and Pritsker, the empirical paper by Kambhu estimates the positive feedback trading that could have occurred in response to a large move in US interest rates in April 1995. Kambhu uses data on derivative dealers' US dollar interest rate options portfolios from the April 1995 Central Bank Survey of Derivatives Markets, market surveys taken by the International Swaps and Derivatives Association, and historical and current data on interest rates and implied volatilities to estimate the risk profile of dealers' US dollar interest rate option portfolios. Chart 9 shows his estimate of the risk profile. The vertical axis of Chart 9 represents the value of dealers' US dollar interest rate option portfolios, and the horizontal axis represents the change in interest rates. The inverted-U shape of the risk profile indicates a negative gamma position,¹² which if dynamically hedged would lead to positive feedback trading.

¹² A "negative gamma position" is one that makes a profit for small changes (up or down) in the underlying market risk factor and a loss for large changes.

Chart 9

The risk profile of dealers' aggregate US dollar interest rate option portfolio



Kambhu goes on to estimate the positive feedback trading that would result from dynamic hedging in response to interest rate shocks of various sizes.¹³ He concludes that the dynamic hedging might be small relative to the size of the market for a small change in interest rates (25 basis points), while a larger change (75 basis points) could produce dynamic hedging that might have an impact only on the medium-term segment of the yield curve, between 5 and 10 years. This research could be extended to other markets and other time periods, perhaps using data from the upcoming 1998 Central Bank Survey. If it were found that negative gamma positions with respect to a particular market risk factor were small, one would have less concern about positive feedback trading, and there would be no identifiable benefit to measuring aggregate market risk for that risk factor.

VI. Concluding remarks

This paper has described the cooperative research effort that was undertaken to explore the measurement of aggregate market risk. The group's overall assessment of the research can be summarised as follows. First, it was felt that the research has not established an adequate technical basis or adequate justification for collecting data on aggregate market risk. It is not clear whether data generated from the books of core intermediaries would convey much information about likely market

¹³ Chart 9 also shows that dealers, in aggregate, had substantial net positive market value in their interest rate option portfolios in April 1995, and it would have taken a large move in interest rates to erase that positive value.

dynamics in the aftermath of a shock, given the large fraction of wealth held - and trading accounted for - by institutional investors. While the group believes that the burden of producing useful aggregate market information will inevitably fall as new information system technologies are deployed, the cost of producing such information would still have to be evaluated relative to the benefits the information would provide.

Second, there was agreement that the research effort had usefully highlighted the importance of the roles played by dealers and non-dealers in market functioning. The financial innovations that have broadened the scope of financial intermediation to include intermediation of price risks have been a positive development that might be expected to lower risk premia in asset prices. These forms of intermediation often rely on the ability of dealers to manage their risks actively. Such risk management requires access to trading and market liquidity, and in their absence dealers would exact higher premia for their intermediation services. Since a broad range of economic agents appear to depend on the liquidity of the core trading markets, either directly or indirectly, more research on market dynamics, market liquidity and market functioning in times of stress would be useful.