Risk measurement and systemic risk: a summary

1.   Overview

Research on risk measurement and systemic risk-related issues, the focus of the conference, has progressed substantially since 1995, when the first in this series of conferences was held. At the first conference, centre stage was taken by the value-at-risk (VaR) methodology, which was then gaining ground in academia and at leading financial institutions. Some papers explored how risk could be quantitatively measured and what the meaning of such measures would be. Shortly thereafter, in 1997, the Asian crisis erupted, triggered by and itself triggering events that were beyond the bounds envisioned by standard VaR methodology. As a result, discussions at the second conference in 1998 very much focused on the implications of the Asian crisis for risk measurement methodologies as well as market microstructure theory's lessons for market dynamics in times of stress.

In his opening remarks, Andrew Crockett explained the rationale for the focus of this third conference and its emphasis on questions relating to the nature and sources of market liquidity, recent advances in risk measurement methods, sources of banking crises and contagion effects across regions and markets. As for the first two conferences in the series, the goal was to foster the exchange between the policy and research communities. To this end, the co-organisers brought together a broad mix of attendees: academics, public sector officials and industry professionals as well as central bank staff. Overall, the conference generated a set of interesting discussions which sought to both assess and further the current state of knowledge on issues related to risk measurement and systemic risk and to identify areas of policy interest and for future research. These discussions focused on three broad topics, which are summarised below under three headings.

2.   Systemic banking crises, contagion and monitoring

The series of banking and currency crises that emerged in various parts of the world during the past two decades or so suggests that financial stability is not to be taken for granted. In view of this, the conference organisers encouraged submission of research concerned, among other things, with the sources of financial market instabilities and, by extension, ways to avoid financial crises. Much of this literature has focused on issues of banking stability and the notion of "systemic risk", ie the danger that problems in a single financial institution might spread and, in extreme situations, such contagion could disrupt the normal functioning of the entire financial system.

Banking stability and systemic crises

Diamond and Dybvig,[1] in their seminal paper, present a theory of banking based on liquidity risk sharing, with banks emerging as providers of the required liquidity insurance. They show how, under asymmetric information, bank runs can emerge in such a fractional reserve banking system. However, while allowing for the possibility of bank runs, the Diamond/Dybvig (DD) model is not able to explain the causes of banking crises: bank runs, in their world, are essentially self-fulfilling prophecies or "sunspot" events.

Extensions of the DD model, as surveyed by Allen and Gale's contribution to this proceedings volume, have therefore introduced uncertainty about asset returns to proxy for the impact of the business cycle on the valuation of bank assets. In these models with aggregate shocks to asset returns, financial crises are driven by fundamentals. Shocks to asset returns, by reducing the value of bank assets, raise the possibility of banks being unable to service their commitments. Depositors, anticipating such difficulty, will tend to withdraw their funds early, possibly precipitating a crisis.

---

[1]   D Diamond and P Dybvig, "Bank runs, deposit insurance, and liquidity", *Journal of Political Economy* 91, 1983, pp 401-19.

Despite its widespread use in theoretically analysing financial instability, the DD model and its various extensions do not provide a completely plausible description of actual patterns of banking crises. Runs by depositors are rare. Therefore, banking crises have more typically started when the interbank supply of credit was sharply cut or withdrawn. In addition, a purely bank-centric approach to systemic risk may no longer be appropriate, given that financial markets tend to play a significant role as propagation channels for disturbances involving the banking system and the real economy. This is why Yutaka Yamaguchi, in his luncheon address, set out the need for any comprehensive analysis of systemic risk to go beyond the narrow confines of the banking system, to cover the interrelations between the banking system, financial markets and the real economy. Indeed, one of the recurring themes of the conference was that much of the literature on banking crises and contagion, the topics of the first two conference sessions, remained overly focused on a set of specific assumptions and modelling conventions. As a result, while being more tractable, these models have provided only limited analytical assistance to the policy community.

In the latest version of their 1998 model,[2] the main focus of the first presentation at the joint research conference, Allen and Gale introduce a market for long-term assets into the analysis, enabling banks to liquidate these assets. Contrary to the original DD model, liquidation costs are therefore endogenous. As a result, asset markets provide a transmission mechanism that serves to channel the effect from the liquidation of assets by some banks to other banks in the economy. If a sufficient number of banks are forced to liquidate their assets and the demand for liquidity rises above a certain level, asset prices will move sharply. This may, in turn, force other banks into insolvency and exacerbate the original crisis. As a result, the model, compared with earlier theories, provides a more realistic explanation of how and why financial crises may develop. It also highlights the importance of asset market liquidity for the evolution and, eventually, the avoidance of financial crises.

Carletti et al, in their presentation, tackled another major shortcoming of many analyses based on the traditional Diamond/Dybvig approach: the failure to recognise the role of interbank credit. In their model, banks compete in the loan market, while the interbank market serves as an insurance mechanism against deposit withdrawals due to liquidity shocks. This setup enables the authors to investigate the influence of bank mergers on reserve holdings and the interbank market and, ultimately, aggregate liquidity risk. Mergers affect bank balance sheets via increased concentration and potentially enhanced cost efficiency, while also altering the structure of liquidity shocks. The model highlights the importance of functioning interbank markets for financial stability and sheds some light on potential trade-offs between antitrust and supervisory policies. In the discussion, some conference participants commented on the practical relevance of the model. In particular, it was noted that nowadays central banks were usually ready to provide liquidity elastically to accommodate temporary fluctuations in liquidity. Given this willingness, it was argued, bank liquidity crises would be of limited importance. However, it was felt that the paper generated important insights into how mergers might affect liquidity in the money market and, by extension, how this would influence the execution of monetary policy operations.

The final presentation of the first conference session, which is summarised in Giannetti's contribution to this volume, shifted the focus to the emerging markets. Specifically, she argued that underdeveloped financial markets, characterised by a lack of transparency, and easy access to foreign capital can help to explain overlending and crisis phenomena in emerging financial markets. According to Giannetti, overlending due to investor moral hazard, that is the existence of explicit or implicit guarantees, is merely a special case of a broader crisis model. In her model, based on incomplete investor information on the average quality of investment opportunities and the existence of soft budget constraints due to capital inflows, bank-financed investors will rationally not require a risk premium until losses become substantial, even without guarantees on deposits. Based on this insight, the paper suggests that well developed capital markets, by increasing the number of creditors, can eliminate excessive reliance on bank-firm relationships and soft budget constraints, which will reduce the probability of financial crises. This, in turn, lends support to the often advocated "sequencing" policy prescription, demanding that countries should have appropriate financial structures in place before removing capital controls and passively accommodating foreign investors.

---

[2] F Allen and D Gale, "Optimal financial crises", *Journal of Finance* 53, 1998, pp 1245-84.

Contagion

During the second session, the conference's focus moved on to contagion across markets and countries, an issue which, despite its importance for financial market stability, remains less than completely understood. Contagion is at the heart of any analysis of financial crises, because it is contagion that makes the initial shock a truly systemic event. Therefore, echoing Tommaso Padoa-Schioppa's luncheon speech, to understand financial sector risks, one has to deal with the origin of these risks as well as the channels of propagation. Padoa-Schioppa noted that the increasing use of complex risk transfer instruments and speed of financial market transactions add to the complexity and rapidity of the potential propagation of shocks, making these risks difficult to gauge. Based on these considerations, contagion can be viewed as the propagation mechanism that causes small idiosyncratic or systematic shocks to have systemic consequences.[3]

The session started with two empirical presentations, by Kaminsky and Reinhart and Dungey et al, investigating contagion by using stock market and bond market data, respectively. Kaminsky and Reinhart's presentation involved an analysis of daily stock market behaviour for a number of emerging and mature markets. Specifically, their paper looks at empirical return distributions in different countries and regions, conditional on extreme returns in financial centres or emerging markets, to identify where shocks originate and how they spread through the system. By comparing these distributions, the authors discover that the distribution of returns around the globe changes only in those periods that are characterised by turmoil in large financial centres (notably the United States, Germany and Japan). While shocks might spill over regionally, via trade links, centres have to be affected for financial turmoil to be become a global phenomenon. That is, shocks to the periphery seem to spread to other peripheral areas via their impact on financial centres. A shock that never reaches a centre is likely not to become a systemic event.

Mardi Dungey and her co-authors employ a somewhat different approach. They identify contagion by looking at daily movements in bond spreads for the LTCM crisis period in an effort to quantify the effects of unanticipated regional shocks across borders. The resulting contagion measure controls for common global shocks, country specific shocks and regional factors. The authors find contagion originating from the Russian default, with the measured level of the effect larger for emerging economies. However, the proportion of total volatility attributable to contagion varies widely across countries and is not always more substantial for developing countries. Thus, while contagion tends to be viewed as mainly a concern for developing countries, the evidence from the Russian and LTCM crises suggests this is not necessarily the case. In fact, contagion effects are found to be widely distributed across both developed and developing markets, making contagion a phenomenon reserved not only for developing countries.

In the discussion of the two empirical contagion papers, however, doubts were expressed about whether the data and methodologies used in these and similar empirical models were always suitable for identifying the effects of contagion. It was pointed out, for example, that Kaminsky and Reinhart's decision to look at daily stock price returns of emerging economies to establish how turmoil in an emerging market spills over to other markets was open to criticism. In particular, it was observed that the definition of what a crisis is and when it started might change when equity prices rather than exchange rate data are used. For example, during the Asian crisis, days of crisis in stock and foreign exchange markets tended to differ, as stock markets in Asia welcomed the initial depreciation of local exchange rates as a necessary adjustment. As a result, turmoil in their paper might be very different from what is commonly perceived as a crisis, limiting the value and comparability of their findings. In addition, doubts were voiced as to whether the paper actually addressed the issue of contagion, given its focus on patterns of spillovers in stock markets. In particular, conference participants suggested that future work might look more closely at causalities by trying to infer the direction of spillovers. Finally, it was proposed to apply the two papers' methodologies to recent cases of limited or non-contagion. This was seen as potentially useful in testing the hypothesis that, during recent episodes, investors have been more discriminating in their reactions than in the past.

In the third presentation of the contagion session, Cipriani and Guarino elaborated on social learning and informational herding as a source of financial crises. While in the previous presentations, such as

---

[3] See O DeBandt and P Hartmann, "What is systemic risk today?", in *Bank of Japan, Risk Measurement and Systemic Risk: Proceedings of the Second Joint Central Bank Research Conference*, Tokyo, 1999, pp 37-84.

the one by Allen and Gale, financial crises and contagion were essentially based on developments in fundamentals or sunspot phenomena, Cipriani and Guarino advance reasons for crises in the absence of sunspots and despite sound fundamentals. Essentially, the authors introduce the possibility that crisis phenomena might reflect a learning process between traders, independent of any change in fundamentals. By doing so, the approach provides a possible underpinning for the centre-periphery results found by Kaminsky and Reinhart, as increased effects on peripheral markets could now be interpreted as the result of a higher trading frequency at the centre. The authors also implicitly challenge insights from other areas of research, such as second-generation speculative attack models.

According to Cipriani and Guarino, a possible explanation of why sound fundamentals may not be reflected in asset prices is that information about these fundamentals may be spread among investors, with prices failing to fully aggregate it. In particular, this would happen if investors, instead of acting according to their own private information, simply decided to follow the actions of previous traders, a phenomenon known as informational herding. Specifically, the authors use an information cascades model with flexible prices to show that sequential trading under incomplete information can lead to a permanent deviation of prices from fundamentals. In such a model, prices may fail to aggregate private information and may, due to asymmetric information, lead to all traders taking the same action. Under specific conditions, traders choose to essentially disregard their own private information, not allowing asset prices to reflect fundamental values. Furthermore, in a multi-market sequential trading framework, it can be shown that sell orders in one market can affect the price path of another market, making its price settle at lower value. While such informational spillovers are to be expected, due to correlation between fundamentals, sequential trading can explain contagion across markets as correlation between the prices of two assets can be higher than correlation between fundamentals. Informational herding effects can, therefore, spill over from one asset market to the other, providing a potential explanation for contagion across markets.

Conference participants, however, raised doubts about the validity of the two core assumptions behind the model - the existence of "gains from trade" and the restriction that trades occur only sequentially. In particular, it was noted that gains from trade implied that market participants would be willing to trade at a loss. With regard to future research, it was suggested that empirical implementations of the cascades approach could shed some light on contagion effects, eg during the Asian crisis. For this to be possible, however, cascade models would have to be reworked to generate verifiable theoretical predictions on, for example, conditions under which informational cascades were more likely to occur.

Systemic monitoring

Systemic events can impose substantial social costs on the affected economies, as bank runs, for example, will disrupt credit relations and allocative efficiency, in turn leading to non-trivial direct and indirect effects on economic performance in the form of output losses. Practical aspects of systemic monitoring and the analysis of systemic risks are, therefore, high on the policy agendas of central banks and other members of the regulatory and supervisory community. For this reason, the last conference session featured two papers that added a practical angle to the discussion on banking crises and contagion, by showing how financial market and banking data can be used to monitor the fragility of real-world banking sectors.

Against this background, Gropp et al explore how market-based indicators can be usefully employed to predict banking fragility by adding to the information gained from more traditional, balance sheet-based indicators. To this end, the authors analyse the indicator characteristics of Merton-type distances-to-default and subordinated bond spreads in signalling material weakening of banks' financial conditions. They demonstrate that useful and well-behaved indicators can be derived from stock market data, while, so far, the focus has been much on subordinated debt spreads. They also find that these market-based indicators, with different leads, are useful in predicting banking fragility and that they even add information relative to more traditional indicators based on balance sheet information. The authors thus suggest the use of market-based indicators in supervisors' early warning models, a potentially promising future enhancement of supervisors' ongoing monitoring efforts.

Blåvarg and Nimander, in their paper, give valuable insights into the Riksbank's monitoring of systemic risk in the Swedish banking system. In particular, to monitor counterparty exposures in the domestic interbank market, the Riksbank uses data detailing the largest uncollateralised exposures of the four major players in the Swedish banking system. The approach involves exposing a proxy for the Swedish banking system, ie the four biggest banks, to solvency shocks originating from outside the

interbank market and assessing how the system is affected via interbank exposures. The authors find that domestic direct contagion effects are less than what might have been expected in the Swedish banking system, given its degree of concentration. In most cases where one of the large banks is assumed to fail, other banks are found not to suffer direct losses that would reduce their capital ratio significantly below the regulatory level. Similar results are found for the risk of direct contagion from abroad, which mainly arises from foreign exchange settlement exposures. Conference participants suggested that the approach presented might be extended to explicitly take into account correlated shocks due to common exposures. This, and coverage of possible second-round effects of a given primary shock, were avenues suggested for future research.

3. Market liquidity

As argued above, much of what was discussed during the first two sessions revolved around the concept of market liquidity and its relevance for financial stability. Banks, the epicentre of instability in the models surveyed by Allen and Gale, are providers of insurance for liquidity risk. They serve this function by following a liquidity immunisation strategy, implemented via individual asset markets and interbank credit markets, to guard themselves against the possible effects of forced asset liquidation. This, in turn, explains the organisers' motivation for placing particular emphasis on papers seeking to explicate the specifics of liquidity provision in various microstructural seetings and across various asset markets.

Against this background, the papers presented in the two liquidity sessions all addressed issues of liquidity provision and, in various ways, all supported the view that market liquidity can affect market performance, while, in turn, being affected by market microstructure. In the first liquidity session, the paper by Cohen and Shin explored the short-run variability of US Treasury note prices using order flow data from the US Treasury market. The paper by Tien investigated the determination of exchange rates using currency futures data disaggregated by type of trader. Finally, Pritsker employed a theoretical asset pricing model to demonstrate the possibility that the asset holdings of large investors might matter for asset price determination. In the second liquidity session, Danielsson and Payne examined the microstructural specifics of liquidity provision on an electronic foreign exchange trading platform, while Harrison and Wong and Fung looked into the microstructure of the primary corporate bond and the equity markets, respectively.

Positive feedback in the Treasury market

Cohen and Shin explore the empirical relevance of strategic interaction among market participants. In particular, they are interested in whether the distributions of returns/liquidation values are more dispersed than they would be if risks were truly exogenous. A direct implication of such a finding would be that individual market participants are likely to underestimate potential price movements resulting from shocks to markets and, therefore, predictably underestimate the riskiness of their own exposures. The empirical part of the paper investigates return and order dynamics in the US Treasury bond markets to find that signed order flow has a strong impact on prices. While this is fully in line with what one would have expected based on standard market micostructure models, the authors also find that there is often also a strong effect in the other direction, ie prices affecting order flow. As this is found to be more likely in turbulent times, bond markets seem to behave in meaningfully different ways depending on market conditions. This effect in the price-order flow pattern, so the authors claim, may be attributed to constraints on traders' behaviour, such as those imposed by risk management systems or position limits. As these and similar constraints can give rise to "strategic complementarities", the most basic concept of strategic interaction, the actions of individual traders may become mutually reinforcing, introducing feedback from prices to order flow. The specific issue of how VaR constraints might affect asset prices and volatilities, an interesting topic against the background of Cohen and Shin's findings, was taken up again by Berkelaar et al and is covered below.

Market microstructure and FX market liquidity

Tien, in his paper, shifted the attention to the foreign exchange (FX) market and investigated FX premia based on hedge demand, where risk (forward) premia are driven by income shocks and risk-averse agents' attempts to hedge these shocks by trading foreign currency. The model is tested using data on hedging demand in currency futures markets and the author finds evidence indicating that FX risk premia based on hedge demand explain, on average, some 45% of the variation in currency returns at a monthly horizon. Therefore, risk premia are present and identifiable in the foreign exchange market and, more importantly, risk sharing can explain a significant proportion of the observed variation in exchange rates. This, in turn, suggests that the FX market is an efficient mechanism for allocating risk across the economy. These results may also help to explain the depth and liquidity of the major currency markets, since traders should be more willing to trade in situations where counterparties are not likely to be better informed. In the discussion of the paper, it was suggested that the definition of hedgers used in disaggregating the data by type of trader might be a problem that could potentially skew the results. It was separately noted that, by extension, Tien's findings also supported opposition to recent proposals for the introduction of Tobin taxes in the foreign exchange markets. This is because such a tax would interfere with the needs of those market participants seeking to hedge their income risk - an unnecessary burden from a viewpoint of allocative efficiency.

Daníelsson and Payne, in their presentation, remain in the realm of the FX market by empirically investigating liquidity provision on electronic FX broking systems. Such electronic trading platforms, having captured a sizeable market share in the inter-dealer FX market, have recently attracted considerable interest, particularly as they rely on electronic order books and, thus, on limit orders as the ultimate source of liquidity.[4] A deep limit order market would, therefore, be characterised as having a large volume of differently priced buy and sell limit orders outstanding, waiting to be "hit" by market orders arriving in the market. Such a market would thus be able to absorb large numbers of limit orders without significant price movements, while being able to restore the depth of the order book once a market order is executed.

In their latest paper, which was the basis of their presentation, Daníelsson and Payne use DEM/USD Reuters data for a particular week in October 1997 to look at the dynamics of market liquidity. In particular, the authors seek to establish the conditions driving liquidity supply and demand in the market. They find that market order activity has strong and persistent effects on subsequent limit order activity in electronic order books. In addition, they show that the order book is dynamically illiquid in the sense that, subsequent to market order arrival, further liquidity is removed from the other side of the order book as buy orders cause liquidity suppliers to reprice limit orders, leading to a reduction in sell side depth. The order book, therefore, "thins out" as liquidity suppliers seem to guard themselves against being picked off by traders with superior information, a finding in line with market microstructure models based on asymmetric information. In addition, the authors find depth to be negatively related to volatility and unexpected volume, while being positively related to expected volume. This, in turn, suggests that liquidity suppliers are risk-averse and concerned about informed trades by market order traders. On the other hand, as remarked in the subsequent discussion at the conference, such correlation between volatility and market depth could also be a reflection of liquidity providers not being anxious to enter a market where there is not a sufficient background level of volatility to justify their presence. Overall, this line of research was seen as an important contribution to the existing empirical market microstructure literature. Nevertheless, it was noted that there was much more research to be done before arriving at an informed understanding of liquidity generation in order-driven markets.

Market microstructure and stock market liquidity

The paper by Wong and Fung looks into the liquidity of equity markets, using a unique set of 30-second tick-by-tick data from the Hong Kong Stock Exchange. Various conventional liquidity

---

[4] See Committee on the Global Financial System, *Structural aspects of market liquidity from a financial stability perspective: a discussion note prepared by the CGFS for the March 2001 meeting of the Financial Stability Forum (FSF)*, Bank for International Settlements, Basel, 2001.

indicators are constructed to evaluate how liquidity has evolved since the 1997 Asian crisis and to examine the determinants of changes in liquidity. The analysis shows that, having deteriorated during the Asian and Russian financial crises, market liquidity has broadly recovered to pre-crisis levels. In addition, to more fully gauge the dynamics of market liquidity, a GARCH model is developed for five selected stocks to relate the sensitivity of their price movements to net order flows. Based on this model, the authors establish that market liquidity deteriorated sharply during the crises, followed by an apparent recovery in the post-crisis period.

Overall, given the correlation of stock market liquidity with cost and risk factors established in the paper, the authors find their results to be consistent with market microstructure models based on inventory control, predicting that market depth is negatively correlated with price volatility. In such a model, as limit orders are essentially options that can be exercised by submitting a market order, heightened volatility would imply an increased risk for the limit order provider to deviate from his optimal inventory position - which would, in turn, lead to declining order book depth. However, as in Danielsson and Payne, it was noted that a negative correlation between depth and volatility could well be consistent also with microstructure models based on asymmetric information and risk-averse liquidity providers interacting with informed traders. Finally, interest was expressed in further studies to fully reconcile the various theories on how market microstructure might affect market prices with real-world market structures and transactions data from different markets.

Issue size and bond market liquidity

Issue size is known to be an important determinant of bond market liquidity and the issuer's funding cost. For this reason, an empirical study by Harrison investigated the issue size-liquidity linkage by looking into the impact of liquidity shocks on the composition of firms entering the corporate bond market. As much previous research on bond market liquidity has focused on secondary markets, examining the primary bond market provides additional insights into what issue and issuer characteristics may be fundamental liquidity factors. To this end, Harrison's approach focuses on the role of issue size and its sensitivity to illiquidity. That is, unlike other authors, he looks at the effect of market stress on liquidity, rather than the causes of market stress and illiquidity. Using multivariate regressions to control for observable issue and issuer characteristics, he finds that issue size, and certain measures of issuer familiarity, are priced liquidity factors. Primary markets, therefore, seem to recognise and price information problems and related factors of liquidity determination at issuance. In particular, the price depends crucially on whether the economy is experiencing an illiquidity shock. When liquidity is at a premium, larger bonds by well known issuers are much more prominent, squeezing issues by smaller, less known firms and the high-yield market in particular. Overall, it seems, with multiple issues and large issues being discounted, that the prospect of wider ownership translates into more trading and more liquidity for these securities.

In the discussion, there was agreement that, while the paper was more or less agnostic about what exactly explains the link between size and liquidity, it would be worthwhile to examine the issue further. In particular, it was felt that size might well proxy for some very specific factor not (yet) captured in the paper. In addition, it was suggested that the hedging of corporate bond inventories might influence the econometric results, if not properly controlled for. If inventory hedging becomes more expensive, dealers will become more reluctant to bring new issues to the market, making hedging a core factor behind the activity in the primary market. From a policy viewpoint, Harrison's findings were seen to suggest that, as conditions in the primary bond market tend to reflect conditions in the secondary market, those monitoring liquidity can also turn to the primary market to gauge liquidity developments. For example, as the composition of issuers tends to change rather dramatically in response to periods of illiquidity, it may be interesting to look at who is coming to the market instead of just looking at the overall amount issued.

Large investors and market liquidity

Pritsker, to address questions related to liquidity determination, constructs an imperfect competition model of asset pricing without focusing on a particular market. A key innovation of his approach, when compared to the studies surveyed in his contribution to this volume, is the assumption that institutional investors incorporate the price impact of their actions into their own decision-making. Imperfect competition and the existence of agents with differently sized endowments, ie asset holdings, imply that large agents face costs, due to illiquidity, when trying to rebalance portfolios. As a result, large

traders will be hesitant to trade away from their endowments. That is, in response to acquiring an appreciation of the possible consequences of their actions on market outcomes, large traders will tend to sell less of their endowments when subject to liquidity shocks. In turn, observed market returns on assets will be directly related to the size of large traders' endowments. The model thus endogenously generates trading costs and explores their implications for asset pricing and market liquidity. Therefore, this line of research is likely to be useful in further exploring issues such as why some shocks are more contagious than others, or why some assets are more liquid than others.

One conjecture that would follow from this analysis, as remarked in Peter Praet's discussion of the Pritsker paper, is that large traders will not only sell less of their endowments but will also be biased towards holding liquid, blue-chip stocks in their portfolios. It was also noted that, given that Pritsker suggests that large market players may have an incentive to hide their asset endowments, extending his model to an environment with asymmetric information concerning investors' holdings could yield interesting new results. Such an approach could, for example, be used to compare the "full disclosure" case adopted in Pritsker's model with situations of zero disclosure - an interesting undertaking from a policy perspective.

4. Practical risk measurement and management

Over the course of the conference, several participants noted that, with regard to practical risk measurement, substantial progress had been made since the first conference in 1995. There was agreement, however, that further improvements were necessary in terms of modelling the tails of return distributions, improving the treatment of liquidity risk, and integrating the measurement of market and counterparty credit risk. The performance of risk measurement systems in times of stress and possible shortcomings of conventional methods in dealing with such situations received particular attention throughout the discussions. Against this background, conference participants commented on the need for the use of other techniques, such as stress testing, to address the shortcomings of the more traditional risk measures, an issue that also received the attention of two recent CGFS reports.[5]

The papers in the fifth, the technical, conference session applied cutting-edge statistical techniques to specific issues of financial risk measurement. One paper, by Diebold et al, showed how high-frequency data can be employed to construct volatility forecasts which, in turn, can be used as an input for firms' risk measurement. The authors integrate high-frequency intraday FX data into the measurement and modelling of daily and lower-frequency volatility and return distributions, overcoming the problems of more restrictive, traditional approaches in terms of dealing with intraday frequencies. The relevance of the study stems from the fact that volatility forecasting is a prominent feature of many practical financial decisions such as asset allocation, market timing and derivatives pricing.

The second paper, by Yamai and Yoshiba, compared two popular summary measures of financial risk, value-at-risk (VaR) and expected shortfall, using extreme value theory (EVT). The authors use simulated asset returns with extreme correlations and fat-tailed distributions to compare the performance of the two measures under market stress and evaluate whether the measures take account of extreme losses in the tail of the underlying distributions (tail risk) and whether they can be accurately measured using limited data (estimation error). In the open discussion, while agreeing with some of the advantages of expected shortfall measures under conditions of stress, some participants raised doubts as to whether expected shortfall could be a practical measure to be actually used by banks. In particular, it was noted that, while VaR had very good statistical properties, not much was known about the distribution of expected shortfall and that using expected shortfall for backtesting might pose problems.

Lucas et al, in their presentation, used EVT to describe how the tail of the loss distribution in portfolio credit risk models depends on modelling assumptions and parameter choices. While tail index and

---

[5]   See Committee on the Global Financial System, *Stress testing by large financial institutions: current practice and aggregation issues*, Basel, 2000, and *A survey of stress tests and current practice at major financial institutions*, Basel, 2001.

quantile estimators, like VaR, are now commonly used to assess the tails of return distributions, application of these statistical techniques in calculating extreme credit loss quantiles is less common. In their paper, the authors investigate whether the application of extreme value theory to the tails of portfolio credit losses generates EVT quantiles that are accurate enough to be useful for credit risk managers. To this end, alternative tail approximations are considered for two special cases of a generalised model for portfolio credit losses. The results suggest that one has to be careful in applying EVT for computing extreme quantiles efficiently. The applicability of EVT in characterising the tail shape appears to depend crucially on the exact distributional assumptions for the systematic and idiosyncratic credit risk factors. These factors are seen to limit the applicability of standard EVT methods in the credit risk context. This leads the authors to suggest that more care should be taken when using EVT for credit risk management purposes. In the discussion, this last implication of the paper triggered some controversy, as EVT is already widely used throughout the banking sector to model various types of financial risk, including credit risk. With regard to future work, it was suggested that the authors could consider extending their current, one-factor approach to a multi-factor setting to enhance the applicability of their research.

Finally, as part of the sixth and last conference session, Berkelaar et al investigated how the application of standardised, VaR-based risk management tools might reduce market participants' risk-taking in normal circumstances at the expense of increasing exposure to extreme events. Their paper, therefore, sheds light on how practical risk management using now-standard statistical techniques might affect market dynamics and equilibrium prices. To this end, the authors extend earlier research to find that, in a world with VaR-constrained agents, market volatility (as well as implied options volatility) is generally reduced, generating a stabilising effect for the economy as a whole for most states of the world. However, in extremely bad states, agents have an incentive to gamble by taking large exposures, pushing up market risk and creating a hump in the equilibrium price function. As a consequence, losses for most states are thus reduced at the expense of the remaining states where, with the probability of extreme losses fixed via the VaR constraint, losses will be larger than in the unconstrained case. While this was seen as an interesting and potentially important insight, in the course of the discussion conference participants noted that the results were based on highly restrictive assumptions, such as the strict application of VaR limits and the absence of other risk-related constraints. Given these assumptions, the model was seen as being based on an overly rigid notion of risk management. This raised doubts about the direct practical relevance of this particular model's insights. It was, hence, left for future research to investigate the topic further.