Research Summaries

Introduction

Bank of Canada staff undertake research designed to improve overall knowledge and understanding of the Canadian and international financial systems. This work is often pursued from a broad system-wide perspective that emphasizes linkages across the different parts of the financial system (institutions, markets, and clearing and settlement systems), linkages between the Canadian financial system and the rest of the economy, and linkages to the international environment, including the international financial system. This section summarizes some of the Bank's recent work.

In the article, "The Impact of Unanticipated Defaults in Canada's Large Value Transfer System," Darcey McVanel examines how robust individual participants in Canada's Large Value Transfer System (LVTS) are to defaults by other participants in the system. The LVTS is designed to meet international risk-proofing standards at a minimum cost to participants in terms of collateral requirements. It does so, partly through collateralized risk-sharing arrangements, whereby participants may incur losses if another participant defaults. The LVTS is designed to be robust to default. Its rules, however, do not mean that individual participants are robust to default. The author studies participants' robustness to default by simulating unanticipated defaults, calculating the loss allocations that other participants would have to bear, and comparing these loss allocations with participants' collateral in the LVTS and with their capital positions. She finds that all participants are able to withstand the loss allocations that result from the largest defaults that she can simulate using actual LVTS data.

Many countries prohibit large shareholdings in their domestic banks. In "Ownership Concentration and Competition in Banking Markets," Alexandra Lai and Raphael Solomon ask whether such prohibitions hinder competition. The authors study a loans market with two banks. Managers choose loan levels and appropriate part of the cash flow; either a controlling shareholder or the manager chooses the bank's debt. The holders of large blocks of shares (blockholders) are more likely to win control. The authors show that banks with controlling blockholders would issue more debt, since the blockholder "disciplines" the manager by reducing free cash flow. More debt leads the manager to issue more loans, thus providing a more competitive market. Since controlling blockholders result in increased competition, shareholding restrictions inhibit competition. The authors ignore possible self-dealing by blockholders, but note that good governance and banking supervision can address self-dealing. The authors conclude that prohibitions on concentrated ownership merit further study.

Central bankers have a long-standing interest in how financial assets move together over time and, in particular, during times of market stress. To understand this, central bankers need a model of the time-varying covariance matrix of asset returns. In "Using High-Frequency Data to Model Volatility Dynamics," Gregory H. Bauer presents a new model of the covariance matrix that he developed with Keith Vorkink of MIT. The model has several advantages over existing methods. High-frequency data are used to construct daily estimates of the volatilities of, and correlations between, stocks with different market capitalizations. A new mathematical technique is then used to model the evolution of this matrix over time. The authors show that this evolution can be explained by a small number of variables. In the future, they hope to use the model to understand the dynamics of international assets.

The Impact of Unanticipated Defaults in Canada's Large Value Transfer System

*Darcey McVanel**

Canada's Large Value Transfer System (LVTS) is designed to meet international risk-proofing standards at a minimum cost to participants in terms of collateral requirements.[1] It does so partly through collateralized risk-sharing arrangements whereby participants may incur losses if another participant defaults, but the system itself is robust to default. The LVTS is designed so that participants pledge sufficient collateral to cover at least the largest possible payment obligation to the system. This does not mean, however, that *individual participants* are robust to default. Participants are responsible for managing their own risks to protect themselves from potential losses stemming from the default of another participant.[2] In the paper summarized here, the ability of participants to withstand such defaults is assessed by simulating unanticipated defaults in the LVTS. (In reality, there have not been any defaults in the LVTS.)

Key Features

The LVTS forms the core of the Canadian payments system. It substantially reduces systemic risk and allows Canada to meet the best international practices for handling large-value payments by applying the following risk-control elements:

- The net amount that each participant is permitted to owe is subject to bilateral and multilateral limits. Individual payments are subject to risk controls to ensure that they do not exceed these limits.
- At the beginning of each business day, participants pledge collateral to the Bank of Canada with a value sufficient to cover the largest permitted net debit position from a single participant. This will provide the liquidity required to settle the system should one of the participants default.
- The Bank of Canada guarantees settlement in the extremely unlikely event that more than one participant defaults on a single day and that the sum of the exposures exceeds participants' pre-pledged collateral.

These elements provide participants with certainty of settlement for those payments that pass the risk-control tests.

Participants can send their payments through one of two payment streams. In the first stream, participants pledge their own collateral to cover their obligations. This stream is referred to as "defaulter pays," since, in the case of a default, the defaulter's own collateral is used to generate liquidity to settle the system. The second stream is termed "survivors pay," since, in the case of a default, the non-defaulting participants share the costs of settling the defaulter's obligations. While participants in this stream clearly bear risks related to the exposures of other participants, this stream has much lower collateral costs than the first.

In the survivors-pay stream, participants determine the limits of the exposure they are willing to assume vis-à-vis other participants and extend lines of credit accordingly. Each participant must then pledge collateral to cover a standard percentage (currently set at 24 per cent) of the largest bilateral credit limit (BCL) it has extended to any other participant. This is the maximum amount that the participant will have to contribute if one or more participants to which it has granted a BCL defaults. On the reciprocal side, each participant can incur a net bilateral position equal to the BCL that has been established for it by the grantor and a net multilateral

1. For a full description of the LVTS, see Dingle (1998).
2. A participant is in default if it cannot meet its end-of-day net debit position.

* This article summarizes a recently published Bank of Canada working paper (McVanel 2005).

position equal to a fixed percentage of the *sum* of the credit lines granted to it. (See Box 1 for an example.)[3]

Participants who end the day with an overall net debit position must find either the funds or the collateral to settle their position; otherwise, the participant is in default.[4] Since participants in the survivors-pay stream can incur a net debit position that exceeds their collateral, default is possible in the LVTS.

If a participant defaults, its own collateral will first be used to absorb its losses. Other participants will then share in the remaining losses in proportion to the size of the BCLs they have granted to the defaulter. Participants have control over the size of the BCLs that they grant to the defaulter. They also have the incentive to set them small enough to be able, from a solvency perspective, to withstand the losses incurred in the event of another participant's default. In this study, maximum-impact defaults are generated based on actual LVTS data in order to test whether participants are indeed setting BCLs at a level sufficient to withstand their losses.

Box 1

Example of Credit Limits

Participant *A* grants a *BCL* of 10 to participant *B* and one of 20 to participant *C*.

A must therefore pledge collateral of 0.24 (20).

B and *C* grant *BCL*s to *A* equivalent to the *BCL* granted to them by *A*.

A can incur a net debit position of:

- up to 10 with *B*
- up to 20 with *C*
- overall (*B*+*C*) up to 0.24 (10+20) = 7

(Note that, since there are 15 participants in the LVTS, the multilateral constraint is less restrictive than this example would suggest.)

Methodology and Data

The study period spans the 170 business days from 1 March to 29 October 2004. The average daily volume and value of payments over this period were 17,063 and $130.2 billion, respectively. Data on participant transactions, collateral, and bilateral credit limits are used to determine participants' maximum positions, shortfalls, and loss allocations.[5] Participants' Tier 1 capital is used to determine whether they can withstand their losses.[6]

If a participant is closed by its regulator during the LVTS day, it will immediately become ineligible for further participation in the system. Our defaults are generated by assuming that each

3. For a more detailed discussion of credit limits in the LVTS, see McPhail and Senger (2002, 46).
4. Participants can use both the collateral supporting their defaulter-pays obligations, as well as their survivors-pay collateral.
5. We thank the Canadian Payments Association for providing these data.
6. Data for federally regulated financial institutions are obtained from the website of the Office of the Superintendent of Financial Institutions, and data for all others from the websites of the institutions themselves.

Chart 1 Size Distribution of Participants' Losses

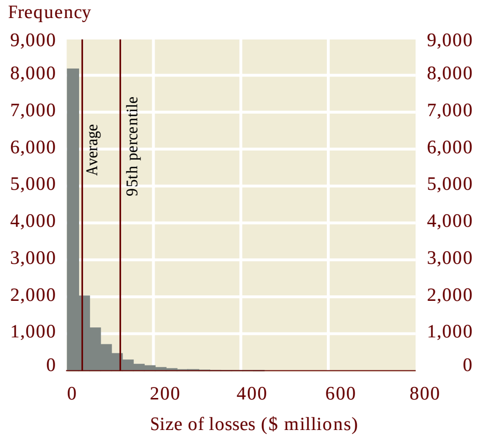

participant is closed by its regulator and is, therefore, ineligible to participate after the point when it reaches its maximum net debit position. Participants' maximum negative positions are found by simulating actual LVTS activity over our time period, using the Bank of Finland Payment and Settlement Simulator.[7] In each case, this position is compared with the participant's collateral to determine whether survivors would incur losses. Survivors' losses are then calculated according to LVTS Rules, with survivors sharing in the losses in proportion to the size of the bilateral credit limit that they granted to the defaulter.[8] Survivors' losses are compared with participants' Tier 1 capital holdings, and participants are deemed able to withstand their loss if their Tier 1 capital after the loss exceeds the level required by their regulator.

Results

A participant is said to have incurred a shortfall in each case where it is closed with a net debit position that exceeds the value of its collateral. Shortfalls occur in almost half of all cases. The size of the average shortfall is relatively small, about 20 per cent of the maximum allowed (based on BCLs granted), and on each participant's worst day, shortfalls are, on average, about 80 per cent of the maximum possible.

Chart 1 illustrates the size distribution of survivors' loss allocations, which are generally very small. Large participants bear nominal losses that are approximately four times larger than those of small participants, implying that the largest losses are borne by those participants most able to bear them. Loss allocations as a proportion of Tier 1 capital are very small—just 0.35 per cent, on average. But small participants absorb the largest loss allocations as a proportion of Tier 1 capital, especially on the worst days, meaning that small participants take on relatively more risk. In the worst case, losses can be as high as one-third of capital. Even here, however, the participant's capital remains higher than that required by its supervisor. Therefore, even the most significant loss would not cause any participant to fail.

7. We thank the Bank of Finland for providing the Bank of Finland Payment and Settlement Simulator for our use.
8. See McVanel (2005) for the exact formula.

To summarize, LVTS participants are in general easily able to withstand losses resulting from the default of another participant. Furthermore, the losses found in this study are probably larger than would be seen if a participant were actually to default. First, the largest possible shortfalls were created, based on the data, to maximize survivors' losses. Second, the default was assumed to be unanticipated. This prevents participants from reducing or eliminating BCLs to the defaulter to avoid sharing losses. Finally, it was assumed that survivors do not recover any of their losses.

References

Dingle, J. 1998. "The LVTS—Canada's Large Value Transfer System." *Bank of Canada Review* (Autumn): 39–55.

McPhail, K. and D. Senger. 2002. "The Impact of Participant Outages in Canada's Large Value Transfer System." Bank of Canada *Financial System Review* (December): 45–48.

McVanel, D. 2005. "The Impact of Unanticipated Defaults in Canada's Large Value Transfer System." Bank of Canada Working Paper No. 05-25.

Ownership Concentration and Competition in Banking Markets

*Alexandra Lai and Raphael Solomon**

Do restrictions on the ownership structure of banks limit competition? This question is relevant to more than 50 countries, including Canada, that either prohibit individuals and corporations from holding more than a given fraction of a bank's shares or require that large shareholders be reviewed by the government or by the central bank.[1]

While there are good prudential or governance rationales behind rules requiring dispersed shareholdings, these rules have their own drawbacks. For example, they may reduce access to cheaper capital and increase operating costs. This article focuses on the operational problems associated with shareholding restrictions. These problems arise in situations of potential conflict of interest between the different stakeholders of a firm. In this study, we model the conflict of interest that arises between bank shareholders and bank management, and ask whether restrictions on the ownership structure of banks can restrict competition. Since our work is not calibrated to the data of any particular country, and since we model only one potential cost to shareholding rules without modelling their benefits, we cannot directly evaluate any particular country's shareholding rule. We do, however, shed light on a potential cost of shareholding rules that might prove substantial for countries with less than perfectly competitive banking sectors.

There is a substantial literature on ownership concentration. While most empirical work in this area has examined non-bank firms, Caprio, Laeven, and Levine (2004) provide empirical evidence of a positive relationship between ownership concentration and value for a sample of 244 publicly traded banks across 44 countries. There is some evidence of a positive relationship between control by blockholders (the owners of large blocks of shares) and firm performance in the United States. Barclay and Holderness (1989) and subsequent studies confirm that large blocks of shares trade at a premium, evidence of net private benefits from large block ownership. There is also some evidence that block formation and block trades are associated with excess stock price increases, suggesting shared benefits from such control (Mikkelson and Regassa 1991; Barclay and Holderness 1991, 1992). Hence, private benefits need not reduce the wealth of minority shareholders. Indeed, Holderness and Sheehan (1998) present evidence from the United States that large blockholders are constrained from expropriating cash flows and from other actions inimical to the interests of minority shareholders. Barclay and Holderness (1991) further find that this increase in firm value is limited if the blockholder does not exercise control (which they define to be actions such as changing the composition of the board or replacing the management).

All of the above studies deal with blocks held by external investors and not with managerial (inside) shareholdings. Morck, Shleifer, and Vishny (1988) find that firm value initially increases with small amounts of managerial ownership, decreases with managerial ownership for an intermediate range of shareholdings, and then increases again for very large managerial shareholdings. McConnell and Servaes (1990), on the other hand, find that firm value increases with managerial ownership up to 40 to 50 per cent and decreases thereafter.

Key Model Features

To formalize the operational problems associated with shareholding restrictions, we set up a game-theoretic model of two competing banks, in which bank managers choose the level of

1. In Canada, neither individuals nor corporations may hold more than 20 per cent of the voting stock of banks with assets greater than $5 billion.

* This article summarizes Lai and Solomon (2006).

loan activity (quantities) and appropriate a fraction of the bank's residual cash flow for themselves (for example, in the consumption of benefits or perks). But either the bank manager or the controlling blockholder can choose the level of the bank's risky borrowing and, thus, the bank's capital structure.[2] To obtain control, the holder of a block of shares must engage in costly monitoring. Monitoring does not guarantee control, but it gives the blockholder the possibility of control. The more shares the blockholder owns, the more likely it is to win control. If there is no blockholder, or if the blockholder fails to obtain control, then the manager chooses the bank's capital structure. The timing of the game is as follows. First, the two potential blockholders simultaneously decide whether to acquire a controlling share of the bank and whether to monitor management. Next, either the manager or the controlling blockholder chooses the capital structure of their bank. The proceeds of any debt sold are distributed to equity holders, rather than being used to finance operations. Finally, the managers of the two banks compete in the market for loans, repay debt holders, and appropriate residual cash flow.

Results

There are three possible outcomes for the banking industry: (i) both banks are controlled by a blockholder, (ii) both banks are controlled by a manager, or (iii) one bank is controlled by a blockholder and the other by the manager. We find that controlling managers always issue less debt than controlling blockholders. As a result of their debt choices, banks controlled by managers extend fewer loans than those controlled by blockholders. Competition for loans is thus fiercest in an industry where both banks are controlled by blockholders and tamest in an industry where both banks are controlled by managers.

From a blockholder's perspective, issuing debt has two consequences. First, it "disciplines" a manager by reducing the amount of free cash flow from which the manager can appropriate. Second, it creates a strategic effect in the loans market vis-à-vis the other bank, as demonstrated by Brander and Lewis (1986). Specifically, holding fixed the amount of debt at the rival bank, a unilateral increase in one bank's debt induces that bank to extend more loans while inducing the other bank to extend fewer loans.[3]

Why would a manager whose bank has already increased the riskiness of its balance sheet by issuing debt become even more aggressive and expand the bank's loan portfolio? The key is that the bank has limited liability. In the presence of debt, extremely negative shocks give the bank a return of zero, while beneficial shocks give the bank a positive return, which actually increases as more loans are issued. Thus, the issuance of debt by one bank causes that bank's manager to compete more aggressively in the loans market relative to a market where neither bank issues debt. This raises the market share and profits of the indebted bank at the expense of the rival bank, since the issuance of debt makes the industry less profitable overall.

In a symmetric (Nash) equilibrium, where both banks issue debt, each bank's lending operations are less profitable than they would be were the two banks to function as a single (merged) entity. However, an increase in debt at both banks may increase the value of both banks. The commitment to repay debt implicitly transfers resources from the manager to the shareholders. Free cash flow has two uses: repayment of the debt and appropriation for the manager's private benefit. Larger debt repayments necessarily entail less appropriation, thus increasing the value of the bank. Moreover, the banking industry is more competitive than it would be if less debt were issued, and consumer welfare also increases as more debt is issued. Since managers issue less debt than blockholders, the presence of controlling blockholdings increases the value of banks, as well as competition in the loans market.

We find that a minimum size of shareholding is necessary to induce a blockholder to monitor. This is because the probability of winning control and, hence, the expected benefits of control, increase with the size of the block held, while the cost of monitoring is fixed. We also find that this minimum holding is larger for the blockholder facing a rival bank with its own blockholder

2. We do not consider other regulatory constraints, such as minimum capital requirements, that banks face when making portfolio decisions.

3. This is a simple result of downward-sloping reaction functions arising from the Cournot game.

than it is for the blockholder facing a rival bank with dispersed ownership.[4]

We distinguish three classes of bank shareholding rules that restrict ownership concentration to a designated level: (i) non-restrictive—the maximum shareholding is such that a blockholder would monitor management even if the rival bank also had a blockholder, (ii) moderately restrictive—the maximum shareholding is such that a blockholder would monitor management if the rival bank did not have a blockholder but would not monitor if the rival bank had a blockholder, and (iii) highly restrictive—the maximum shareholding is such that a blockholder would never monitor management, regardless of the ownership structure of the rival bank.

When shareholding rules are non-restrictive, blockholders that subsequently monitor management form at both banks. When shareholding rules are moderately restrictive, blockholders form at both banks, but neither monitors management; hence, industry outcomes are the same as if both banks were widely held. Finally, when shareholding rules are highly restrictive, investors are dissuaded from acquiring blockholdings, and both banks have dispersed ownership.

Implications

Our analysis suggests that legal restrictions on the concentration of ownership can affect the value of bank shares, as well as competition in the loans market. Shareholding restrictions affect banking competition through the capital structure of the bank. Our model does not, however, consider regulatory capital requirements that may affect the decisions of either blockholders or managers regarding capital structure. Marginally relaxing the shareholding restriction will affect competition only in cases where the restriction has not prevented blockholding and monitoring from occurring. If ownership restrictions are severe enough to prevent blockholding or monitoring (even if blockholdings form), then a marginal increase in the maximum shareholding will, generally, not affect bank value or competition in the loans market. For a relaxation of restrictions on bank shareholding to be beneficial, the increase in maximum shareholding may need to be substantial.

Our model also abstracts from other conflicts of interest between equity holders and debt holders (risk shifting) and between blockholders and minority shareholders (self-dealing). While the problem of risk shifting is particularly relevant to highly leveraged institutions, such as banks, capital requirements and positive franchise values mitigate the problem. Moreover, risk shifting is associated with leverage and not with ownership concentration.

Restrictions on bank shareholding date back to the 1960s in some countries. There have since been two important developments. First, corporate governance in the general corporate sector and in the banking sector improved significantly in the 1980s and 1990s. This included changes such as an increased emphasis on outside directors, new rules for electing boards, and more internal oversight. Second, since the implementation of Basel I in 1992, the supervision of banks has increased, particularly that of large, multinational banks. Taken together, these changes vastly reduce the scope for self-dealing by the holders of large blocks of shares. The prevention of self-dealing as a justification for limited concentration, while fairly valid in the 1960s, is, therefore, less important today in most industrialized countries. We believe that it is relevant to consider the potential costs of this regulation, and we have modelled one such cost.

In almost all of our simulations, a rule restricting ownership concentration to no more than 20 per cent leads to two outcomes.[5] In the first, blockholders never exist; in the second, blockholders exist but do not monitor and never gain control. Since we do not calibrate the model (this would require good estimates of the demand for loans, agency costs, and monitoring costs), it is difficult to say whether restricting ownership to 20 per cent is excessive. But our results indicate that restrictions on bank shareholding can discourage monitoring, thus reducing competitiveness in the loans market.

4. This is the case for almost all of the parameterizations in our numerical examples.

5. The median and modal restriction among countries in the World Bank database (Barth, Caprio, and Levine 2001) is 20 per cent.

References

Barclay, M.J. and C.G. Holderness. 1989. "Private Benefits from Control of Public Corporations." *Journal of Financial Economics* 25: 317–95.

———. 1991. "Negotiated Block Trades and Corporate Control." *Journal of Finance* 46(3): 861–78.

———. 1992. "The Law and Large-Block Trades." *Journal of Law and Economics* 35(2): 265–94.

Barth, J., G. Caprio, and R. Levine. 2001. "The Regulation and Supervision of Banks around the World: A New Database." World Bank Policy Research Working Paper No. 2588.

Brander, J. and T. Lewis. 1986. "Oligopoly and Financial Structure: The Limited Liability Effect." *American Economic Review* 76: 956–70.

Caprio, G., L. Laeven, and R. Levine. 2004. "Governance and Bank Valuation." World Bank Policy Research Working Paper No. 3202.

Holderness, C.G. and D.P. Sheehan. 1998. "Constraints on Large-Block Shareholders." National Bureau of Economic Research Working Paper No. 6765.

Lai, A. and R. Solomon. 2006. "Ownership Concentration and Competition in Banking Markets." Bank of Canada Working Paper No. 2006-7.

McConnell, J.J. and H. Servaes. 1990. "Additional Evidence on Equity Ownership and Corporate Value." *Journal of Financial Economics* 27: 595–612.

Mikkelson, W.H. and H. Regassa. 1991. "Premiums Paid in Block Transactions." *Managerial and Decision Economics* 12(6): 511–17.

Morck, R., A. Shleifer, and R.W. Vishny. 1988. "Management Ownership and Market Valuation: An Empirical Analysis." *Journal of Financial Economics* 20: 293–315.

Using High-Frequency Data to Model Volatility Dynamics

Gregory H. Bauer*

The covariance matrix of asset returns is important for a wide range of individuals.[1] Academics use estimates of the covariance matrix to test asset-pricing theories. Portfolio managers use the covariance matrix in designing tracking strategies where the return on their portfolio is designed to closely follow the return on a benchmark portfolio. Risk managers use the matrix to construct measures such as "value at risk." Corporate managers require accurate measures of covariances for hedging strategies.

Central bankers also have a profound interest in this concept. An assessment of financial market stability and contagion depends on measuring the time-varying variances and covariances that make up the matrix. For example, research has shown that there is an "excess" comovement of international equity markets during market downturns (e.g., Connolly and Wang 2003; Ribeiro and Veronesi 2002). Whether this is a rational response to current economic conditions or the result of irrational "contagion" remains an open question.

It is a key stylized fact in empirical finance that the variances and covariances of asset returns fluctuate over time.[2] Central bankers and others, therefore, require a model of a time-varying or "conditional" covariance matrix.[3] Several distinct methods for estimating a conditional covariance matrix have evolved in the literature, but since an asset's true volatility cannot be observed, researchers must treat the elements of the covariance matrix as non-observed or "latent" processes. This greatly complicates the modelling of the covariance matrix. If the actual matrix could be observed, the causes of time-varying market volatilities and correlations could be measured more accurately.

Realized Volatility

The concept of "realized volatility" has recently been developed to provide more precise estimates of the volatility of a single asset or index. Assets such as stocks and bonds trade second by second throughout the day. These high-frequency data can be recorded and aggregated to yield a relatively precise estimate of the daily volatility of the asset. The resulting realized volatility is not latent, but observed, which results in more accurate forecasts.[4] While most papers have focused on estimates of the volatility of a single asset, it would be interesting to see whether a better estimator of the entire conditional covariance matrix could be created in this way.

In "Multivariate Realized Stock Market Volatility," Gregory Bauer (Bank of Canada) and Keith Vorkink (MIT) introduce a new model of the conditional covariance matrix. High-frequency data for a number of stocks are recorded during

1. A *covariance* measures how the price of one asset moves over time in relation to the price of another. A *covariance matrix* is a mathematical concept that measures how several asset prices move together over time. It is composed of the variances of the individual assets and the covariances between them.
2. For a comprehensive survey of the literature on volatility modelling and forecasting, see Andersen et al. (2005).
3. "Conditional" refers to market participants using current information to make optimal forecasts.

* This summary is based on Bauer and Vorkink (2006).

4. Andersen and Bollerslev (1998) introduced the idea of using high-frequency data to construct estimates of the daily realized volatility of a single asset. Andersen et al. (2003) formalized the definition, which was applied to equity markets in Andersen, Bollerslev, Diebold, and Ebens (2001) and exchange rates in Andersen, Bollerslev, Diebold, and Labys (2001). Constructing realized volatilities requires care because of the institutional trading features present in high-frequency data.

the day. Once aggregated, the data can be combined to construct estimates of the daily conditional covariance matrix. By using this approach, the variances and covariances of a number of assets can be treated as being observed. As a result, more accurate estimates of the factors driving the conditional covariance matrix can be found.

Bauer and Vorkink apply their new approach to the cross-section of size-sorted U.S. stock portfolios. While earlier papers have examined asset-price volatility in the cross-section of small and large firms,[5] they used existing models of latent volatility to capture the variation in the covariances. In contrast, Bauer and Vorkink use high-frequency data to construct daily measures of the realized covariance matrix of small and large firm return indexes over the 1988 to 2002 period. Their measures of volatility are more precise than those in previous work and allow for a more detailed examination of the causes of conditional covariances.

Once the matrix of realized variances and covariances has been constructed, a new factor model is used to capture its dynamics.[6] The factors are functions of past volatilities and other variables that can help forecast future volatility. A number of possible sets of variables from the academic finance literature are then examined to see how well they forecast the covariance matrix. The authors note that while researchers have examined different variables for their ability to forecast stock market returns, there is much less evidence that the variables forecast stock market volatility.[7]

Results

Bauer and Vorkink evaluate their model of the daily conditional covariance matrix in two ways. First, they use a set of standard statistical tests and find that, in general, the factor model performs well in describing how the volatility matrix changes each day. Surprisingly, however, there does not appear to be a lot of difference between the alternative forecasting variables used to construct the factors: one set of variables appears to forecast the covariance matrix just as well as another set. This is because a single dominant factor drives the volatilities of all of the different-sized stocks: if the overall market is volatile, then the prices of all stocks on that day are volatile. As long as the forecasting variables are able to capture the dynamics of aggregate market volatility, they will also capture the dynamics of the size-sorted stocks.

The second and more informative method of evaluating the model is to see how well it constructs optimal stock market portfolios. In particular, the authors examine how the model can be used to construct a daily "tracking-error" portfolio.[8] The covariance matrix of the size-sorted stocks is modelled, and the indexes are used to track the portfolio of "value" stocks (i.e., those with high book-to-market ratios). Including variables that forecast stock returns (such as dividend yields) along with lagged volatility factors is found to produce portfolios with superior tracking performance. In other words, variables that forecast returns also forecast risks (i.e., volatility) in the market.

The authors hope to use this method to explore the time-varying relationship among other asset markets and to determine how well alternative variables are able to forecast large movements in market prices. The model can also be used to examine the covariances among international assets with a view to better understanding the transmission of shocks from one country to another, especially during times of market stress.

5. See Conrad, Gultekin, and Kaul (1991); Kroner and Ng (1998); Chan, Karceski, and Lakonishok (1999); and Moskowitz (2003).
6. In the factor model, the variances and covariances of a large number of assets are explained by a small number of variables.
7. For example, there is evidence that a stock market's dividend yield (the dividend-to-price ratio of the index) may help predict the average return on the index, but whether it predicts the volatility of returns (from holding the index) is unknown.
8. A tracking-error portfolio is one in which the portfolio manager uses a small set of assets to "track" or closely follow the performance of the target portfolio. The idea is to minimize the difference between the returns on the tracking and target portfolios. For example, fund managers may combine a number of stocks and derivative products to match the performance of a broad equity market index, such as the TSX composite index. The manager may thus trade in only a few assets to follow the returns on many stocks, which would greatly reduce transactions costs. Because the tracking-error portfolio test is based on the difference between the volatilities on the tracking and target portfolios, it is less influenced by moves in aggregate market volatility that affect both portfolios.

References

Andersen, T.G. and T. Bollerslev. 1998. "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts." *International Economic Review* 39: 885–905.

Andersen, T.G., T. Bollerslev, P.F. Christoffersen, and F.X. Diebold. 2005. "Volatility Forecasting." National Bureau of Economic Research Working Paper No.11188.

Andersen, T.G., T. Bollerslev, F.X. Diebold, and H. Ebens. 2001. "The Distribution of Realized Stock Return Volatility."*Journal of Financial Economics* 61: 43–76.

Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys. 2001. "The Distribution of Realized Exchange Rate Volatility." *Journal of the American Statistical Association* 96: 42–55.

———. 2003. "Modeling and Forecasting Realized Volatility." *Econometrica* 71: 579–625.

Bauer, G.H. and K. Vorkink. 2006. "Multivariate Realized Stock Market Volatility." Bank of Canada. Photocopy.

Chan, L.K.C., J. Karceski, and J. Lakonishok. 1999. "On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model." *The Review of Financial Studies* 12: 937–74.

Connolly, R.A. and F.A. Wang. 2003. "International Equity Market Co-Movements: Economic Fundamentals or Contagion?" *Pacific-Basin Finance Journal* 11: 23–43.

Conrad, J., M.N. Gultekin, and G. Kaul. 1991. "Asymmetric Predictability of Conditional Variances." *The Review of Financial Studies* 4: 597–622.

Kroner, K.F. and V.K. Ng. 1998. "Modeling Asymmetric Comovements of Asset Returns." *The Review of Financial Studies* 11: 817–44.

Moskowitz, T.J. 2003. "An Analysis of Covariance Risk and Pricing Anomalies." *The Review of Financial Studies* 16: 417–57.

Ribeiro, R. and P. Veronesi. 2002. "The Excess Comovement of International Stock Markets in Bad Times: A Rational Expectations Equilibrium Model." Graduate School of Business, University of Chicago Working Paper.