

A framework for assessing the systemic risk of major financial institutions[†]

Xin Huang^{a,*}, Hao Zhou^b and Haibin Zhu^c

^a*Department of Economics, University of Oklahoma*

^b*Risk Analysis Section, Federal Reserve Board*

^c*Bank for International Settlements*

This version: May 2009

Abstract

In this paper we propose a framework for measuring and stress testing the systemic risk of a group of major financial institutions. The systemic risk is measured by the price of insurance against financial distress, which is based on *ex ante* measures of default probabilities of individual banks and forecasted asset return correlations. Importantly, using realized correlations estimated from high-frequency equity return data can significantly improve the accuracy of forecasted correlations. Our stress testing methodology, using an integrated micro-macro model, takes into account dynamic linkages between the health of major US banks and macro-financial conditions. Our results suggest that the theoretical insurance premium that would be charged to protect against losses that equal or exceed 15% of total liabilities of 12 major US financial firms stood at \$110 billion in March 2008 and had a projected upper bound of \$250 billion in July 2008.

JEL classification: G21 ; G28 ; G14 ; C13

Keywords: Systemic risk; Stress testing; Portfolio credit risk; Credit default swap; High-frequency data

* Corresponding author. Tel.: 405 325 2643; fax: 405 325 2643.

E-mail addresses: xhuang@ou.edu (X. Huang), hao.zhou@frb.gov (H. Zhou) and haibin.zhu@bis.org (H. Zhu).

[†]We would like to thank Mark Carlson, Sean Campbell, Francis X. Diebold, Darrell Duffie, Robert Engle, José L. Fillat, Michael Gibson, Brenda González-Hermosillo, Michael Gordy, Don Nakornthat, Nikola Tarashev, George Tauchen, and seminar participants at the Bank for International Settlements, the Hong Kong Institute for Monetary Research, Federal Reserve Board Finance Forum, Hong Kong University, International Monetary Fund, Federal Reserve Bank of Kansas City, and the conference participants at the Financial Econometrics and Vast Data Conference organized by the Oxford-Man Institute of Quantitative Finance, Federal Reserve System Committee Meeting on Financial Structure and Regulation at Boston, Federal Reserve Bank of Chicago's 45th Annual Conference on Bank Structure and Competition, Shanghai Winter Finance Conference, BIS Second Asia Research Network Workshop on Financial Markets and Institutions, Stress Testing Workshop sponsored by China's Banking Regulatory Commission and International Finance Corporation. We would also like to thank Clara Gacia for excellent data support. A grant by CAREFIN Bocconi Centre for Applied Research in Finance is gratefully acknowledged. The views presented here are solely those of the authors and do not necessarily represent those of the Federal Reserve Board or the Bank for International Settlements.

1. Introduction

Banks have been the most important financial intermediaries in the economy, by providing liquidity transformation and monitoring services. The mal-functioning of the banking system can be extremely costly to the real economy, as illustrated in a number of financial crises in both industrial and developing economies in the past few decades, including the current global credit-liquidity turmoil. Therefore, financial regulators and central banks have devoted much effort to monitoring and regulating the banking industry. Such regulation has been traditionally focused on assuring the soundness of individual banks. More recently, there has been a trend towards focusing on the stability of the banking system as a whole, which is known as the macro-prudential perspective of banking regulation (see Borio (2003, 2006)). For instance, Goodhart et al. (2005, 2006), Goodhart (2006) and Lehar (2005) propose measures of financial fragility that apply at both the individual and aggregate levels. At the international level, the Financial Sector Assessment Program (FSAP), a joint IMF and World Bank effort introduced in May 1999, aims to increase the effectiveness of efforts to promote the soundness of financial systems in their member countries.

In order to assess the health of a financial system, two related questions need to be addressed. First, how to measure the systemic risk of a financial system, where systemic risk defined as multiple simultaneous defaults of large financial institutions? Second, how to assess the vulnerability of the financial system to potential downside risks?

In answering the first question, traditional measures have focused on banks' balance sheet information, such as non-performing loan ratios, earnings and profitability, liquidity and capital adequacy ratios. However, given that balance sheet information is only available on a relatively low-frequency (typically quarterly) basis and often with a significant lag, there have been growing efforts recently to measure the soundness of a financial system based on information from financial markets. For example, Chan-Lau and Gravelle (2005) and Avesani et al. (2006) suggest to treat a banking system as a portfolio and use the *n*th-to-default probability to measure the systemic risk by employing liquid equity market or CDS market data with a modern portfolio credit risk technology. Similarly, Lehar (2005) proposes to measure systemic

risk, defined as the probability of a given number of simultaneous bank defaults, from equity return data. The market-based measures have two major advantages. First, they can be updated in a more timely fashion. Second, they are usually *forward-looking*, in that asset price movements reflect changes in market anticipation on future performance of the underlying entities.

In addressing the second question, stress-testing is a popular risk management tool to evaluate the potential impact of an extreme event on a financial firm or a financial sector.¹ The stress testing exercise typically consists of two major steps. In the first step, an economic model is used to examine the dynamic linkages between the asset quality and underlying driving factors (macro-financial variables or latent factors). In the second step, stress testing scenarios (either historical or hypothetical ones), which are based on extreme movements of the driving factors, are fed into the model to assess the resilience of the financial sector. Avesani et al. (2006) and Basurto and Padilla (2006), among others, are examples of stress testing exercises on the financial sector using market-based information.²

In this paper, we propose a framework for measuring and stress testing the systemic risk of the banking sector. Our framework follows the direction of using market information, but with interesting extensions that are designed to overcome a number of shortcomings in existing studies.³

Echoing some earlier studies, we propose to construct the measure of systemic risk based on *forward-looking* price information of two highly-liquid markets, the credit default swap (CDS) spreads and the equity prices of individual banks. Both are available on a daily basis in *real time*. We are able to derive two key default risk parameters, the (risk-neutral) probability of

¹See CGFS (2000, 2005) and Drehmann (2008a,b) for definitions of stress testing exercises and survey of market practices. Stress testing can be implemented to assess the market risk, as in Alexander and Sheedy (2008), or the credit risk, as in this paper.

²In a recent related research, Hancock and Passmore (2008) propose a vector auto-regression in value-at-risk (VAR in VaR) approach, in which systemic and macroeconomic outlook shocks are first fed into a VAR to compute subordinated debt return movements, and then these debt movements are translated into changes in the bank market value using a Merton-type option pricing model. Finally they construct a VaR measure to compute the amount of capital to protect banks against systemic risks.

³Our methodology is based on publicly available market information. Our framework is not related to the supervisory assessments that were conducted over February-April 2009, which relied on confidential supervisory information to assess potential future losses.

default (PD) of individual banks and the asset return correlations, from the CDS spreads and the co-movement of equity returns, respectively. This approach does not rely on the balance sheet or accounting information that may be available only on a quarterly or longer time frequency, with a significant reporting lag.

Similarly in the stress testing exercise, following the recent studies, we adopt an integrated micro-macro model, which not only examines the impact of general market developments on the performance of individual banks, but at the same time incorporates the feedback effect from the banking system to the rest of the economy. More importantly the joint vector auto-regression (VAR) system employs the financial market variables like market return, market volatility, short rate, and yield spread, that are available at a daily frequency in real time.

Our main contribution is to propose to use a new indicator to assess the systemic risk of the banking sector: the price of insurance against large default losses in the banking sector in the coming 12 weeks. The new measure is economically intuitive, in that it is equivalent to a theoretical premium to a risk-based deposit insurance scheme that guarantees against most severe losses for the banking system. The new measure also has the property that it increases in both PDs and asset return correlations. In other words, an increase in the indicator, or a higher systemic risk, can reflect market participants' perception of higher failure risk as well as their view that the probability of common failings is higher (see Das et al. (2007) and Duffie et al. (2008)).⁴ In addition, the new indicator reflects the various degrees of importance of different banks in contributing to the systemic risk, in that banks are treated heterogeneously based on their relative size.

We also propose a novel approach to estimating the asset return correlation, a key parameter to determine the risk profile of a portfolio. The approach employs an advanced technology in the high-frequency literature, i.e. estimating realized correlation from the intra-day high-frequency co-movements in equity prices.⁵ This technique makes it possible to estimate the

⁴The rather homogeneous sample of twelve large US banks and the turbulent period 2001-08, may be coincide with similar movements of PDs and correlations. However, in an on-going research project focusing on a diversified sample of Asian country banks, we find that the effects of PDs and correlation are quite distinguishable.

⁵Throughout this paper, "realized correlation" is a terminology that refers to correlations calculated from high-frequency intra-day data. This is different from "historical correlation" as calculated from daily data and

asset return correlation in a very short time horizon (e.g. one week). Relatedly, we argue that, to calculate the indicator of systemic risk, a *forward-looking* rather than a historical measure of the asset return correlation is the appropriate default risk parameter to be used. Importantly, we find that realized correlations in the short time horizon provide strong and additional predicting power in forecasting the movement in asset return correlations, relative to equity market and term structure variables.

We apply our approach to 12 major U.S. banks during the sample period 2001-08. We produce a weekly time series of systemic risk indicators, that reflect time-varying market perceptions on the systemic risk of the banking system in the United States. The indicator was stable and at low levels at most times but exhibited substantial increases during market turmoil, e.g. the 2002 credit market deterioration and more remarkably after the inception of the subprime crisis in mid-2007.

Furthermore, the peaks of the systemic risk indicator align well with periods of major adverse developments in the market, such as March 2008. In particular, the systemic risk indicator, the theoretical *insurance premium* required to protect against default losses that equal or exceed 15 % of total liabilities, stood at 110 billion USD in March 2008 and had a projected upper bound of 250 billion USD in July 2008. Remarkably in terms of back testing, in our in-sample quarterly horizon forecasting exercise, the realized systemic risk indicators lie out of the 95% predicted confidence interval in approximately 3.5% (13 out of 375) of sample weeks, which is a strong validation of our integrated micro-macro model.

The remainder of the paper is organized as follows. Section 2 outlines the methodology. Section 3 introduces the data and Section 4 presents empirical results based on an illustrative banking system that consists of twelve major commercial and investment banks in the U.S. financial system. The last section concludes.

2. Methodology

Our framework for assessing and stress testing the systemic risk of a financial system consists of the following major components. First, we estimate two major components that determine the “observed correlation” (ex post observation of correlation).

risk profile of a portfolio, the probability of default and the asset return correlation. Second, we construct an indicator of the systemic risk of a financial system, the price of insurance against large losses of the banking sector, based on the forward-looking PDs and correlations in the next period (a quarter). Third, for stress testing purpose we examine the dynamic linkages between default risk factors and a number of macro-financial factors. An integrated micro-macro model framework enables us to investigate the two-way linkages between the banking sector and the macroeconomy. Lastly, we define stress testing scenarios and explore their implications on the stability of the banking system. Below we explain the methodology in detail.

2.1. Estimating risk-neutral PDs

The PD measure used in this study is derived from single-name CDS spreads. A CDS contract offers protection against default losses of an underlying entity; in return, the protection buyer agrees to make constant periodic premium payments. The CDS market has grown rapidly in recent years, and the CDS spread is considered to be a superior measure of credit risk to bond spreads (see Longstaff et al. (2005), Blanco et al. (2005), Zhu (2006) and Forte and Peña (2009), for example) or loan spreads (see Norden and Wagner (2008)). Following Duffie (1999) and Tarashev and Zhu (2008a), it is straightforward to derive the risk-neutral PD from the observed CDS spread ($s_{i,t}$):

$$PD_{i,t} = \frac{a_t s_{i,t}}{a_t LGD_{i,t} + b_t s_{i,t}} \quad (1)$$

where $a_t \equiv \int_t^{t+T} e^{-r\tau} d\tau$ and $b_t \equiv \int_t^{t+T} \tau e^{-r\tau} d\tau$, LGD is the loss-given-default and r is the risk-free rate. The assumptions required for the above characterization of risk neutral PDs are: constant risk-free term structure, flat default intensity term structure, and recovery risk independent of default risk.

There are three elements in the implied PD estimated from the CDS market: (1) the compensation for actual default losses; (2) default risk premium; (3) other premium components, e.g. liquidity risk premium. Our systemic risk indicator incorporates the combined effects of the above three elements on the price of insurance against distressed losses in the banking

system. Although there is no convincing quantitative framework to decompose these effects, it is generally agreed that the default risk premium and liquidity risk premium explain the majority of the increases in CDS spreads entering the subprime crisis. One piece of evidence is that market estimates of actual default rates (e.g. EDF data provided by Moodys KMV) only increased mildly during our sample period, suggesting the hike in CDS spreads is mainly due to lower risk appetite, or concerns on counterparty risk and liquidity risk premium.⁶

Several remarks are worth noting. First, the PD implied from the CDS spread is a *risk-neutral* measure, i.e. it reflects not only the *actual* default probability but also a risk premium component as well.⁷ This has important implications on the choice of the appropriate indicator of systemic risk. In particular, it might be misleading to use a nth-to-default indicator, as it is typically considered to be a physical measure (see the discussion below).⁸

Second, the PD implied from the CDS market is a *forward-looking* measure, i.e., it reflects the average risk-neutral PD of the underlying entity during the contract period. Hence, it offers a market assessment from a different perspective from what most balance sheet variables (such as bank profitability and non-performing loan ratios) do, which tell what *has* happened rather than what *will* happen for the underlying firm. Under the efficient market hypothesis, market prices should incorporate all relevant information including those from the accounting books, especially for major indices and large firms.

Third, throughout this exercise we adopt the standard assumption of a flat term structure of the default intensity (as reflected in equation 1). This assumption might be violated in reality.⁹ However, some preliminary evidence suggests that this assumption only causes small

⁶This is consistent with recent studies by Tang and Yan (2006) and Bongaerts et al. (2008).

⁷There are extensive studies regarding the difference between the risk-neutral and physical PDs, see Amato and Remolona (2003), Huang and Huang (2003), Eom et al. (2004), Berndt et al. (2005) and Driessen (2005), among others.

⁸If one is more interested in the *physical* measure of systemic risk, or the *actual* probability of bank distress, one should use the physical PD measure (see Berndt et al. (2005) and Lehar (2005)). In contrast, our approach is an internally consistent risk-neutral measure of the systemic financial distress.

⁹In general, if the default intensity has an upward term structure, our assumption will lead to an over-estimation of 1-year PD. On the contrary, if the default intensity has a downward term structure, our assumption will lead to an under-estimation of 1-year PD. Unfortunately, there is no consensus view on which term structure is more appropriate (probably it is an important reason why the constant default intensity assumption becomes a norm because of its simplicity).

bias and will not affect our major results.¹⁰

2.2. Forecasting asset return correlations

Regarding the other key dimension of portfolio credit risk, the default correlation, there exist two popular approaches. One approach estimates it directly from historical data on defaults (Daniels et al. (2005), Jarrow (2001), Das et al. (2007), and Duffie et al. (2008)). However, this approach can lead to substantial estimation errors because defaults are rare events, particularly for portfolios comprising high credit-quality firms, like major U.S. commercial and investment banks. The other approach derives the default correlation indirectly by estimating the underlying asset return correlation from equity or credit market data. The logic behind this approach is that equity (or debt) is a call (or put) option on underlying firm assets. Hence, the comovement in equity prices (or CDS spreads) tends to reflect the comovement among underlying asset values. In practice, Hull and White (2004) propose to use the equity return correlation as a proxy for the asset return correlation,¹¹ the proprietary Global Correlation model by Moody's KMV derives the underlying asset value from equity market data and firms' balance sheet information, and then computes the asset return correlation (see Crosbie (2005)), and Tarashev and Zhu (2008a) derive the asset return correlation from the comovement of CDS spreads.

This paper follows the second approach and adopts the suggestion by Hull and White (2004) to use the equity return correlation as a proxy for the asset return correlation. There are two main reasons. First, equity is the the most liquid type of asset traded in the market. Changes in market conditions and the default risk of an entity will be immediately reflected in its stock price movements. Second, tick-by-tick data are only available in the equity market. The advanced technology in the high-frequency literature makes it possible to compute reliable

¹⁰One possible hint we can get is to compare between 1-year and 5-year CDS spreads of the sample entities. They are highly correlated. At the beginning of our sample period when CDS spreads were generally low, the 1-year CDS spreads were lower than 5-year CDS spreads, implying an upward default intensity curve. However, since late 2007 – when the CDS spreads increased substantially – the 1-year CDS spreads have been more or less in line with 5-year CDS spreads, supporting the constant default intensity assumption. Putting together, the potential bias caused by the constant default intensity assumption tends to have only a small effect on our results.

¹¹See Appendix A for a strict proof and the conditions under which the two correlations are equal.

realized correlation over a very short time horizon (e.g., one week) that has been impossible for daily observations.¹² The short-term realized correlation turns out to add significant predicting power on the future correlation movement.

The logic of using the equity return correlation as a proxy for the asset return correlation lies in the fact that, when the firm leverage is constant, the asset return correlation equals the equity return correlation. When the firm-leverage is time varying, this relationship breaks down and the magnitude of the discrepancy depends on the comovements between asset returns and leverages and comovements between changes in firm leverages as well. And because this condition is more likely to hold approximately true in the short run, we compute equity (asset) return correlations over time horizons that are not longer than one quarter. For instance, using Moody’s KMV estimates of market values of equities and assets, we can calculate time series of leverage (monthly) for each of the 12 banks. We test the hypothesis that leverage is constant over a one-month (two/three/six months) window. The hypothesis is not rejected for 11 (10/7/4) banks. This partly supports our claim that the constant leverage assumption is reasonable in a short time horizon (less than one quarter).

Importantly, we deviate from previous studies by not relying merely on *past* correlation measures, but using *forecasted* asset return correlations to measure portfolio credit risk. This makes our correlation measure consistent with the PD measure, and therefore our indicator of systemic risk will be forward-looking. In forecasting the asset return correlation over the next period (one quarter), we derive the relationship between future realized correlation (ex post observed) and current-period (quarterly and weekly) correlations and a number of other explanatory variables:¹³

$$\rho_{t,t+12} = c + k_1\rho_{t-12,t} + \sum_{i=1}^l k_{2i} \cdot \rho_{t-i,t-i+1} + \eta X_t + \nu_t \quad (2)$$

where ρ refers to the average asset return correlation and the subscript refers to the time

¹²See Appendix B for detailed description on the estimation of realized correlation using high frequency intraday data.

¹³Driessen et al. (2006) derive a market-based, forward-looking correlation measure from the option market. However, option-implied correlations can only be calculated for a portfolio for which both the index and individual entities are actively traded in the option market. The application of their approach is quite limited for the purpose of our exercise of measuring the systemic risk.

horizon (one week as one unit) to calculate the correlations, and X includes a list of financial market variables as detailed in Section 4. Interestingly, we find that short-term (one-week) correlations have significant and additional forecasting power on future (one-quarter) correlations.

2.3. Building an indicator of systemic risk

Once the two key portfolio credit risk parameters are known, we are able to use the portfolio credit risk methodology (see Gibson (2004), Hull and White (2004), and Tarashev and Zhu (2008a)) to come up with an appropriate indicator of the systemic risk for a pre-defined group of banks. In this paper, we propose a “distress insurance premium” – the theoretical price of insurance against financial distress. To compute the indicator, we first construct a hypothetical portfolio that consists of debt instruments issued by member banks, weighted by the liability size of each bank. The indicator of systemic risk is defined as the theoretical insurance premium that protects against distressed losses of this portfolio in the coming 12 weeks. Technically, it is calculated as the risk-neutral expectation of portfolio credit losses that equal or exceed a minimum share of the sector’s total liabilities.¹⁴

We choose this indicator over a few alternative measures, such as the probabilities of joint defaults, credit value-at-risk (VaR) and expected shortfalls (see Avesani et al. (2006), Inui and Kijima (2005), and Yamai and Yoshihara (2005)). One important reason is that our PD measures, which are derived from the pricing of CDS contracts, are risk-neutral. This implies that any indicator constructed based on them is also risk-neutral. However, the above alternative measures are conventionally interpreted as *physical* rather than *risk-neutral*, and hence are more likely to be misinterpreted by users.¹⁵ Even if the researcher is aware of the difference, it is not straightforward to explain to the management how a risk-neutral measure differs from a physical measure from a portfolio perspective. By contrast, our indicator of

¹⁴The premium is represented as per unit of exposure to the hypothetical portfolio, therefore is unaffected by the growing magnitude of the total liabilities in the banking sector. As a complementary measure, we also report the total insurance cost in dollar term in the baseline example (see Section 4.1.2).

¹⁵By contrast, the Lehar index refers to the physical (or actual) probability of joint defaults and is logically intuitive. It is complimentary to our systemic risk indicator; and the combination of these information can shed light on the important question whether the changes in our systemic risk indicator are driven by movements in actual default rates or changes in the risk premium component.

systemic risk has a very intuitive economic interpretation: it is equivalent to the premium for a hypothetical risk-based deposit insurance scheme, which covers all credit losses so long as the loss exceeds a minimum share of the total liabilities of the banking system. Moreover, our indicator has the property that it increases in both PDs and correlations, which is consistent with the general impression that a higher systemic risk is either driven by higher failure rates of individual banks or a higher exposure to the same risk factor.¹⁶ Lastly, the probability of joint default measures treat all banks as equal and do not take into account differential impacts of failures of different size banks.

In calculating this indicator, we rely on Monte Carlo simulations to estimate the unconditional (risk-neutral) probability distribution of portfolio credit losses.¹⁷ We assume that the loss-given-default (LGD), the third dimension of credit risk components, follows a stochastic distribution and is independent of the PD process. In particular, we assume that LGD follows a symmetric triangular distribution with a mean of 0.55 and in the range of $[0.1, 1]$. The mean LGD of 0.55 is taken down from the Basel II IRB formula, which is also consistent with the data. For instance, Markit provides both CDS spreads and the LGD parameters corresponding to each CDS spread. Previous studies (e.g. Tarashev and Zhu (2008a)) show that the average LGD parameter used by market participants in the CDS market is about 60%.¹⁸

2.4. Designing stress testing scenarios

To implement a stress testing exercise, we first need to build the links between the macro-financial part of the economy and the portfolio credit risk parameters, the PDs and correlations. Then we use either history or simulation to examine the impact of shock to the system, and consequently the effects on our systemic risk indicator.

We design an integrated micro-macro model to examine the determinants of PDs and correlations. The *macro* part of the model adopts a VAR framework that allows for dynamic

¹⁶The nth-to-default probability does not have this property, see Section 4 for further discussion.

¹⁷See Tarashev and Zhu (2008b), Appendix B, for the details of the Monte Carlo simulation procedure.

¹⁸The adoption of symmetric triangular distribution follows Tarashev and Zhu (2008b) and is not essential to our results. In a robustness check exercise, we assume that the LGD follows a beta distribution, another popular choice in the credit risk literature, and find little changes in the systemic risk indicator (the results are available upon request). However, using triangular distribution is computationally more efficient, especially for the stress testing exercise (bootstrapping).

linkages between the credit risk factors of the banking system and a list of macro-financial variables that reflect the developments of the macroeconomy and the general financial market. In the VAR analysis, the health of the banking system is affected by the general market conditions, and there is also a feedback effect in the opposite direction. The second (*micro*) part of the model explains the determination of the default risk of individual banks by the credit risk factors of the financial system and other financial market variables. To summarize, the model estimation consists of two parts:

$$X_t = c_1 + \sum_{i=1}^p b_i \cdot X_{t-i} + \epsilon_t \quad (3)$$

$$PD_{i,t} = c_{2i} + a_i \cdot PD_{i,t-1} + \gamma X_t + \mu_{it} \quad (4)$$

where equation (3) represents the macro-perspective of the model, in which X includes the credit risk factors (average PD and one-week correlations) in the banking sector and macro-financial variables. Equation (4) examines the movements of individual PDs in response to changes in market conditions. The results, in combination with the forecasted correlations as estimated from equation (2), form the whole dynamics of the economic system that are relevant for the stress-testing exercise.

Our integrated micro-macro model is different from some existing studies, which rely on analysis of latent factors that drive the comovements of default risk of individual banks (see Avesani et al. (2006), for example). The major disadvantage of the latent factor framework is that the hypothetical scenarios, which are based on the statistical distribution of latent factors, lack a clear economic interpretation. On the other hand, a number of empirical studies, including Amato and Luisi (2006), Ang and Piazzesi (2003), Duffie et al. (2007) and Männasoo and Mayes (2009), have shown that default risk is closely related to the state of the business cycle and the condition of the financial market. Our choice of using observed macro-financial shocks in designing the stress testing scenarios is consistent with the latter approach.

In the final part of the analysis, we design stress testing scenarios based on hypothetical or historical shocks to variables within the VAR system. We feed the shocks into the dynamic

macro-micro system. And the resulting movements in state variables affect the forecasted default risk of individual banks and the forecasts of correlations, which together change our indicator of systemic risk of the banking system.

In designing the stress testing scenarios, we adopt two approaches. The first approach, a purely hypothetical one, specifies the stress testing scenarios based on the statistical properties of the shock variables in the model. In particular, we use the bootstrapping technique to simulate the path of shock terms, including shocks in credit risk factors and macro-financial factors (ϵ , μ , and ν in equations 3, 4 and 2) in the next 12 weeks. For each simulation, the impact of the indicator of systemic risk is re-calculated. The simulation is implemented for a large number of times, and the stress testing scenarios are defined as the set of scenarios that generate the most remarkable increases in systemic risk—the 95% quantiles of the path realizations.

The second approach uses historical scenarios, i.e. shocks that occurred during well-known market turmoil periods. However, given that some data (CDS spreads and intraday equity data) are only available in a recent short period, we have to rely on a smaller VAR model that includes only macro-financial variables but can be estimated in a longer sample period (back to 1986). The smaller VAR system includes all macro-financial factors that are included in equation (3). The shocks in macro-financial factors are then fed into the system and the impact on the systemic risk is examined.¹⁹

The two approaches are complementary and provide a general picture of the vulnerability of the financial system, from both the statistical and the historical perspectives. The first approach is more generally a *forecasting* exercise and the results are richer in terms of describing the possible movements—both improvements and deterioration in systemic risk. The second approach, instead, focuses only on the downside risk, with a major advantage of being easily interpretable and connected to major historical crises.

¹⁹The shocks in credit factors are assumed to be zeros, as there are not actual observations of CDS spreads, although other treatment may be viable for us to experiment in the future.

3. Data

The proposed methodology outlined in Section 2 is general and can apply to any portfolio that consists of entities with publicly tradeable equity and CDS contracts in the market. For illustrative purposes, we analyze the banking system in the United States over the period 2001-2008. Our banking group consists of 12 major banks in the United States, namely Bank of America, Bank of New York, Bear Stearns, Citibank, Goldman Sachs, JP Morgan Chase, Lehman Brothers, Merrill Lynch, Morgan Stanley, State Street Corp, Wachovia and Wells Fargo. They represent the biggest commercial banks and security firms in the U.S. and therefore their portfolio credit risk has a direct and major impact on the health of the U.S. financial system.²⁰

Our sample data cover the period from January 2000 to May 2008. We retrieve weekly CDS spreads from Markit, compute realized correlations from high-frequency intraday equity price data provided by Trade and Quote (TAQ) (see Appendix B for the methodology), and retrieve a list of macro-financial variables that reflect the general condition of the macroeconomy and the financial market (see Appendix C for details). See Table 1 for the summary statistics of credit factor variables and financial market variables.

Given that no restriction has been imposed in the process of estimating realized correlations, the correlation estimate has a general correlation structure, i.e. without a factor-loading structure. Although it does not impose any difficulty in computing the indicator of systemic risk, it is impractical to forecast the correlation structure with all pairwise correlation coefficients to be freely determined. First, our sample data do not provide enough degrees of freedom. Second, it is not guaranteed that the forecasted correlation matrix will be positive definite. For these reasons, throughout this paper we assume the same pairwise correlation coefficient across the correlation matrix.²¹

²⁰Our sample period ends in May 2008. In March 2008, the Federal Reserve facilitated the acquisition of Bear Sterns by JP Morgan Chase. Subsequent acquisitions of Merrill Lynch, by Bank of America, and Wachovia, by Wells Fargo, occurred after the end of our sample period.

²¹The removal of dispersion in pairwise correlation coefficients only has a small effect on the magnitude of the indicator of systemic risk, and almost negligible effects on the dynamics of the indicator. This is probably due to the rather homogeneous banking system examined in this study.

The solid lines in Figure 1 plot the observed time series of our variables of interest. Average risk-neutral PDs, implied from CDS spreads of individual banks and weighted by the size of bank liabilities, peaked in March 2008 toward the end of the sample period. Average correlations, on the other hand, are somewhat higher during both 2002-2003 and 2007-2008.²² The four panels in Figure 2 plot the weekly times series of financial market variables, including the fed funds rate, the term spread (defined as the difference between 10-year and 3-month constant maturity Treasury rates), the one-month return and implied volatility of the S&P 500 index.²³

4. Empirical results

In this section, we report the empirical results for the banking group of interest, using the methodology outlined in Section 2. We first illustrate the calculation of the indicator of systemic risk of the banking system during the sample period, then assess its vulnerabilities to extreme shocks in the stress testing exercise.

4.1. Constructing the indicator of systemic risk

In order to calculate the indicator of systemic risk of a banking system, we need to know the PDs of individual banks and the corresponding asset return correlations in the future period (one quarter). The risk-neutral PDs can be easily derived from the observed CDS spreads (see equation (1)).²⁴ The “future” asset return correlation, however, is not directly observable and has to be estimated.

In all regressions here and below, PD and correlation variables are transformed so that they can be defined in the range of all real numbers. We perform the Logit transformation on

²²The “future” correlation uses the observed realized correlation in the forthcoming quarter. It is an *ex post* measure and therefore cannot be used directly to calculate the indicator of systemic risk *ex ante*.

²³We do not include macroeconomic variables because of their availability only at a lower frequency. In a robustness exercise, we also include a longer list of financial variables, e.g. the whole term structure of Treasury rates. We then use the principal components (see Allenspach and Monnin (2006)) in our model analysis. This modification does not lead to improvement in the performance of our integrated micro-macro model.

²⁴We set the LGD to be 55% in this exercise.

PD (between 0 and 1), i.e.

$$\tilde{PD} = \log\left(\frac{PD}{1 - PD}\right)$$

Similarly, we perform the Fisher transformation for correlation coefficients ρ (between -1 and 1):

$$\tilde{\rho} = \frac{1}{2} \log\left(\frac{1 + \rho}{1 - \rho}\right)$$

4.1.1. Forecasted asset return correlations

Table 2 examines the determinants of future asset return correlations, measured by the equity return correlations observed in the next quarter. We run three regressions to illustrate that estimating realized correlations from the high-frequency data is helpful for the forecasting exercise.²⁵

In the first regression, the explanatory variables only include realized correlations estimated over one-quarter and one-week time horizons. The one-week realized correlations are supposed to incorporate very recent changes in the correlations and therefore are helpful to predict future correlations. This is supported by the regression results: all explanatory variables have significant and positive effects on the correlation in the next quarter, with an R^2 of 0.54. This is quite striking given that the dependent variable and explanatory variables cover non-over-lapping sample periods, indeed, with a lag of 12 weeks.

The second regression excludes the short-term (one-week) realized correlations, and, instead, includes a list of current-period market factors, including the fed fund rate, the term spread, the S&P 500 return and implied volatility of the current quarter. It arguably represents the best effort one can achieve to explain future correlations without resorting to realized correlation measures. It turns out that only lagged one-quarter correlation and current S&P 500 return are significant in explaining future correlations. It is also meaningful that correlation is persistent — high lagged correlation leads to high future correlation, and that low

²⁵A possible justification of using current-period asset return correlation as a proxy for future asset return correlation, as adopted in existing studies, is that correlations may follow a random walk process. This assumption, however, is not supported by the data in this sample, nor in the study by Driessen et al. (2006).

market returns lead to high comovement (high systemic risk). However, this regression is similarly successful to the first one in term of R^2 (0.55).

The third and last regression includes all explanatory variables mentioned above, which reaches an R^2 of 0.56 and maintains the sign and significance of lagged correlations and market returns. The results provide evidence that movements in short-term realized correlations incorporate important and additional information (compared to the macro-financial variables) on the future movements in correlations.

The dash-dotted lines in Figure 1, in the lower two panels, plot in-sample predictions of *future* asset return correlations. Although they are not perfect, they do catch the trend of correlation movements and perform better than alternative estimates. For instance, our predictions (the above third regression) yield a mean squared error of 0.0036, significantly lower than the mean squared error of 0.0051 if the current-period asset return correlation is directly used as a proxy.²⁶

4.1.2. Indicator of systemic risk in the banking sector

Based on individual PDs and forecasted asset return correlations, we compute the indicator of systemic risk, the theoretical price of insurance against distressed losses in the banking sector over the next three months. As an example, we define “distress” as a situation in which at least 15% of total liabilities of the financial system are defaulted.²⁷ Given that PDs of individual banks are risk-neutral, the price of insurance against distress equals the expectation (under the risk-neutral world) of portfolio credit losses that equal or exceed the pre-defined threshold.

For this purpose the latest portfolio credit risk technology is applied. In particular, we rely on the Monte Carlo simulation method as outlined in Tarashev and Zhu (2008b), Appendix B. The simulation method consists of two steps. In the first step, we simulate the joint default scenarios based on the information of individual PDs and the asset return correlation. In the second step, conditional on defaults occurring in the first step, we simulate the realization

²⁶The difference is statistically significant at the 95% level.

²⁷The 15% threshold is empirically chosen for illustrative purpose. We tried alternative threshold values (e.g. 10%, 20% and 30%) and the results are very similar. In general, the choice of threshold values affects the level of systemic risk indicators, but not their trend.

of LGDs and the overall credit losses of the whole portfolio. Notice that this methodology might be computationally burdensome, but has a major advantage that it is very general. In particular, it fits the purpose of our exercise because the portfolio has the following characteristics: (i) PDs of constituent entities are heterogeneous; (ii) The underlying instruments are unequally weighted; and (iii) LGDs are stochastic and independent of PDs.

Figure 3 plots the price of insurance against portfolio credit losses that equal or exceed 15% of total liabilities of the banking system, with the top panel as per unit of overall exposures (i.e. total liabilities) and the lower panel in dollar terms. The indicator started from about 10 basis points in the first half of 2001, increased and reached a peak of about 35 basis points in the second half of 2002, when high corporate defaults were reported. The indicator then trended downward and reached its lowest level in late 2006 and early 2007. Since August 2007, the indicator rose sharply and peaked around March 2008, and dropped dramatically after the Federal Reserve facilitated the acquisition of Bear Sterns by JP Morgan Chase. In dollar terms, the highest theoretical insurance premium was around \$110 billion in March 2008, well exceeding the amount of the Federal Reserve's \$30 billion non-recourse loan to JP Morgan Chase. Perhaps the market was anticipating a larger default, or it reflected the hike in default risk premium during the market turmoil. Notice that the trend follows very closely with the average PD series in the banking system (see Figure 1, the upper panel), but is also substantially affected by the movement in correlations (though to a lesser extent). For instance, the peak of the indicator coincides with the peak in both PDs and correlations. In addition, comparing between early 2001 and early 2003, the indicator is higher in the second period when the correlation is higher but the PD is more or less the same.

The impact of PDs and correlations on the indicators is more rigorously examined in the regressions in Table 3. The regression shows that our indicator of systemic risk, the price of insurance against distressed losses, increases in both PDs and correlations and the coefficients are highly significant.²⁸ This is consistent with the conventional view that higher default

²⁸Quantitatively, a one-standard-deviation increase in average PDs (0.0053) moves up the indicator by 11 basis points, and a one-standard-deviation increase in average correlations (0.0681) increases the indicator by 2 basis points. It suggests that changes in PDs have a dominant effect on the indicator; the correlation impact exists but plays a secondary role. Hence, although using realized correlations can improve the work, we

rates and higher exposures to common factors are both symptoms of higher systemic risk. By contrast, the (risk-neutral) *n*th-to-default probability measure, another indicator used in other studies (such as Avesani et al. (2006)), does not have this property. In fact, an *n*th-to-default measure typically increases in PDs but may decrease in correlations (see Table 3). Therefore, using *n*th-to-default probability measures will produce at best unsatisfactory, sometime misleading, indicators on the systemic risk of the banking system.²⁹

These results are quite intuitive based on the knowledge of models of portfolio credit risk. Essentially, our measure is *similar* to the spread of a senior tranche in a portfolio, when there are only *two* tranches and the other tranche covers the losses up to the pre-specified threshold. By contrast, the *n*th-to-default probability measures often correspond to some mezzanine tranches in a multiple-tranche secularization structure. It is well known that, when correlations increase, the probability of zero default and many defaults increases but the probability of an intermediate number of defaults decreases. Therefore, the impact on the mezzanine tranches, or equivalently the *n*th-to-default probabilities, is ambiguous. By contrast, for a two-tranche structure, an increase in correlations will always lower the spread of the equity tranche and increase the spread of the senior one.

4.2. Stress testing

We first report the regression results of the integrated micro-macro model, and then implement the stress testing exercise.

4.2.1. VAR analysis (the “macro” part)

The “macro” part of the model refers to a VAR analysis, with the endogenous variables consisting of two credit risk factors – average PDs and current-period correlations – and a number of macro-financial factors. The optimal number of lags in the VAR system is chosen

consider a second-best solution in the application of our method is to use simpler correlation estimates. This is particularly important if high-frequency data are not available in the banking system of interest. Indeed, in a robustness check exercise we adopt historical correlations directly to recalculate the systemic risk indicator. The indicators exhibit very similar dynamics although the levels can be different (the results are available upon request.)

²⁹The results of (risk-neutral) probability of joint defaults are not reported here but are available upon request.

by the Schwarz Information Criteria, which equals one period. Table 4 reports the regression results and the dash-dotted lines in Figures 1 and 2 plot the in-sample prediction of endogenous variables.

All endogenous variables are positively serial-correlated. In addition, there is strong evidence of dynamic linkages among the endogenous variables. The average PD is positively and significantly affected by the average correlation and negatively and significantly affected by the return in the market index. The results are intuitive. Higher systemic risk in the form of elevated correlation leads to more default. The deterioration of the general market (lower market returns) increases the probability of defaults.

The average correlations are negatively and significantly affected by the two interest rate variables, the fed fund rate and the term spread. This may suggest that, when the monetary policy is eased, most banks' asset returns move together more closely. By contrast, when the monetary policy is tightened, banks are affected to a different degree depending on their position in liquidity and equity capital, and the composition of their assets and liabilities. As expected, lower market returns are associated with higher correlations, as a phenomenon of the downside risk. Finally, PDs have a positive effect on correlations, as it should be — defaults are usually clustered (Das et al. (2007)).

While the VAR framework allows for a feedback effect from the banking system to the macroeconomy and the general financial market, the evidence of the feedback effect is quite weak during our sample period. One exception is that the average PD in the banking system has a negative effect on federal funds rates, suggesting that the central bank's interest rate policy may be affected by financial stability concerns in practice. Finally, the positive effect of average PD on the VIX index may be consistent with notion that VIX is regarded as the “market gauge of fear” by practitioners.

4.2.2. Determination of PDs of individual banks (the “micro” part)

The “micro” part of the model investigates the determination of individual PDs, as a function of lagged dependent variables and the current period market variables, including the average PD and correlations in the banking system and the list of macro-financial factors. Table 5

summarizes the regression results.

For all banks, the individual PD series are positively serial-correlated. They are also positively and significantly affected by the average PD in the banking system, and half of the banks' PDs are positively and significantly affected by the average correlation. Regarding the macro-financial factors, the impacts are quite heterogeneous across banks. First, they do not always have a significant impact on individual PDs. Second, when a macro-financial factor has a significant impact on the PD of an individual bank, the sign is not always in the same direction (except for the market return variable, which always has a negative impact if significant). For instance, changes in fed fund rates have significantly positive impacts on four banks but significantly negative impacts on three other banks. This may reflect the different business models and the different balance sheets of the sample banks during the period under review.

4.2.3. Stress testing

Based on the above regression results, the stress testing exercise can be implemented in three steps. In the first step, we choose hypothetical stress testing scenarios based on the VAR regression results. In the second step, these hypothetical shocks are fed into the model to derive the future dynamic movements (up to twelve weeks) in all endogenous variables in the VAR framework. By extension they will affect the future movements in risk-neutral PDs of individual banks (regression results in equation 4) and forecasted correlations (equation 2). In the third and last step, we construct future movements in the indicator of systemic risk under the stress test scenario using the predicted PD and correlation measures.

As described in Section 2.4, there are two approaches to designing the stress testing scenarios. The first approach, a statistical one, adopts the bootstrapping technique to simulate the shocks in the next 12 weeks based on the sample regression of model (including equations 2, 3 and 4). The simulation results are reflected in the future movements of the indicator of systemic risk as shown in Figure 4. On average, 12 weeks after May 2008, the indicator is forecasted to move up slightly to about 0.41%, roughly the levels of late 2007 and late 2002. In the worst 2.5 percentile scenarios, the indicator will jump up above 1.11%, roughly the level

of March 2008 peak; and in the best 2.5 percentile scenarios, the indicator will move down to 0.09%.

The second approach uses historical scenarios, i.e. the shocks of macro-financial factors as observed in two historical turmoil periods: the LTCM crisis (which uses the shocks between July 3, 1998 and September 18, 1998) and the September 11 episode (which uses the shocks between August 31, 2001 and November 16, 2001). The results are shown in Figure 5. In both scenarios, the systemic risk indicator of the banking system increases and reaches a level comparable to the January 2008 about 0.4-0.5% of the total liabilities. Overall, both results suggest that the vulnerability of the banking system would rise ranging from moderately to dramatically as of May 2008.

A major advantage of the bootstrapping stress testing exercise is that it can simulate the distribution of future movements in the systemic risk, which can be used as a *forecasting* tool. In Figure 6, we adopt such a bootstrap forecasting approach for the whole sample and plot the mean and 95% confidence interval of the forecasted indicator 12 weeks ahead. It is clear that the predicted mean (dash line) generally tracks well the realized systemic risk indicator. The confidence interval band was wide from 2001 to 2004 with a local high in late 2002 and the price tag then was about 1% of total liability (the upper bound). Then the confidence band gradually narrowed until early 2007, when predicted mean also stayed low. Of course, the mean prediction of the systemic risk indicator shot up since mid-2007 and the confidence interval band widened as well. The uncertainty peaked in March 2008, then dropped significantly (though still at high levels) after the strong intervention of central banks.

The in-sample performance of our model forecasts is extremely good, which is a strong validation of our model for measuring and stress testing the systemic risk. Remarkably, out of the 375 weeks of predictions, 13 weeks (3.5%) have the realized systemic risk indicators lying outside of the 95% predicted confidence interval. The outliers were concentrated in two periods, the inception (late July to mid September in 2007) and the peak (March 2008) of the subprime crisis. The location of realized indicators within (or outside of) the confidence interval bands is an indicator of the surprising component of the evolution of the systemic

risk. For instance, the severity of the rapid deterioration in market conditions during the two outlier periods is well beyond market expectations. By contrast, the government intervention and consequently the recovery of the market since late March 2008, although exceeding the average of market expectations, are not extremely surprising.

5. Conclusions

In this paper we propose a framework for measuring the systemic risk of a group of major financial institutions. The methodology is general and can apply to any pre-selected group of firms with publicly tradeable equity and CDS contracts. Our approach adopts the advanced high-frequency technique in estimating realized correlation, and shows that short-term realized correlation helps to predict future movements in the asset return correlation. An indicator of systemic risk, which is based on *ex ante* measures of risk-neutral PDs and correlations, offers an insight on the market perception on the level of a theoretical insurance premium that protects against distressed losses in the banking system. The indicator is higher when the average failure rate increases or when the exposure to common factors increases.

The paper also proposes a framework for stress testing the stability of the banking system, based on an integrated micro-macro model that takes into account dynamic linkages between the health of the financial system and macro-financial conditions. The combination of historical and statistical scenarios offer a general picture of the vulnerability of the banking system in the next quarter.

Our study is only a first step toward improving our understanding of the relationship between financial stability, monetary policy and the real economy. The macro-prudential view, which calls for a closer monitoring of asset prices and the stability of the financial system, has become more widely accepted. However, its implementation largely depends on the feasibility on the operational side. The methodology described in this paper may provide a useful starting point along that direction.

Appendix

A. Relationship between equity and asset return correlations

In the framework of Merton (1974), suppose that the market value of a firm's underlying assets follows a stochastic process:

$$dV = \mu V dt + \sigma V dW \quad (5)$$

where V is the firm's asset value, μ , σ are the drift term and the volatility of the asset value, W is a Wiener process.

The firm has only two types of liabilities, debt and equity. The debt has a book value of X and is due at time T . Merton shows that the equity value is determined by:

$$E = VN(d_1) - e^{-rT} XN(d_2) \quad (6)$$

where $d_1 \equiv \frac{\log(\frac{V}{X}) + (r + \frac{\sigma^2}{2})T}{\sigma\sqrt{T}}$ and $d_2 = d_1 - \sigma\sqrt{T} = \frac{\log(\frac{V}{X}) + (r - \frac{\sigma^2}{2})T}{\sigma\sqrt{T}}$.

Under the condition that r , σ and $\frac{V}{X}$ are all constant, it is straightforward that the equity value is proportional to the asset value (because both d_1 and d_2 are constant and X is proportional to V). Therefore, $d(\log(E)) = d(\log(V))$, where $d(\cdot)$ represents first difference. The equity return correlation, under this condition, equals the asset return correlation:

$$\text{cor}[d(\log(E_1)), d(\log(E_2))] = \text{cor}[d(\log(V_1)), d(\log(V_2))]$$

B. Estimating equity return correlations from high frequency data

B.1. Data

The raw high-frequency data consist of all the tick-by-tick transaction data for the stocks of the twelve banks traded in the major U.S. stock exchanges, including NYSE, Boston, Philadelphia, Pacific and NASD. The trading time is from 9:30 to 16:00 Eastern Time. The data are subject to market microstructure noise, such as non-synchronized trading and bid/ask spreads. The impact of such noise on the realized correlation depends on our sampling frequency and the market activity. For the twelve major banks studied in this paper, their markets are

quite deep. There are typically more than one trade per second. So following Andersen et al. (2003), we use equally-spaced thirty-minute returns to construct our realized correlation measure. This sampling frequency strikes a balance between mitigating the influence of market microstructure noise and preserving the accuracy of the asymptotic theory underlying the construction of our realized correlation measures.

We use previous tick method to construct thirty-minute price data from the tick data. That is, the last price observation in the previous thirty-minute interval is taken as the price of this thirty-minute mark. Then we compute the thirty-minute geometric returns by taking the difference between two adjacent logarithmic prices.

B.2. Realized Correlation Construction

The vector of the logarithmic prices of the 12 stocks, $p(t)_{12 \times 1}$, is assumed to be a 12-dimension semi-martingale (SM) by the no-arbitrage condition. $t \geq 0$ denotes the continuous time. Then the log price can be written as

$$p(t) = a(t) + m(t)$$

where $a(t)$ is the drift part with finite variation, and $m(t)$ is the diffusion part. Notice that $m(t)$ is a local martingale with possible jump components.

Assume that there are M equally spaced observations for each h time period. In our study, h can be a day, a week or a quarter. Corresponding to our thirty-minute sampling interval, M takes the values of 12, 60 or 8640. Then i 'th period j 'th return is a 12×1 vector, computed as

$$r_{i,j} = p\left(\left(i-1\right)h + \frac{hj}{M}\right) - p\left(\left(i-1\right)h + \frac{h(j-1)}{M}\right), \quad j = 1, 2, \dots, M.$$

The realized correlation coefficient for the i 'th period between stock k and l is

$$\hat{\rho}_{(kl),j} = \frac{\sum_{j=1}^M r_{(k)j,i} r_{(l)j,i}}{\sqrt{\sum_{j=1}^M r_{(k)j,i}^2 \sum_{j=1}^M r_{(l)j,i}^2}} \quad (7)$$

Barndorff-Nielsen and Shephard (2004) proposed the asymptotic theory underlying the above realized correlation measure. In particular, they show that $\hat{\rho}_{(kl),j}$ is consistent for the

unobserved population correlation coefficient $\rho_{(kl),j}$, as the sampling frequency goes to infinity.

$$\hat{\rho}_{(kl),j} \xrightarrow[M \rightarrow \infty]{P} \rho_{(kl),j}$$

Additionally, if the price process is a continuous stochastic volatility semi-martingale, that is, when there is no jump in the price process, then Barndorff-Nielsen and Shephard (2004) show that $\hat{\rho}_{(kl),j}$ is asymptotically conditionally normally distributed.

With the above well-defined asymptotics underlying the $\hat{\rho}$ measure, we compute our realized correlation coefficient measure according to equation (7).

C. Data sources and definitions

Our analysis uses weekly data during the period 2001-2008. The list of variables and their sources are:

1. CDS data are from Markit and include daily CDS spreads for each of the 12 sample banks. The CDS quotes refer to 5-year, senior unsecured, no-restructuring clause and US dollar denomination. We use end-of-week observations to construct weekly CDS data.
2. Realized equity return correlations are calculated from high frequency intraday equity price information of sample banks, using the methodology as described in Appendix B. The tick-by-tick equity data are provided by TAQ. In each week, we calculate the realized correlation measures over different time horizons, from one week to one quarter.
3. Financial variables. They include two variables on the performance of the general financial market, one-quarter return of the S&P 500 index and the implied volatility (VIX) of the index, both of which are available from Bloomberg. In addition, we also include the fed fund rate and the term structure, the latter defined as the difference between 10-year and 3-month constant maturity Treasury rates. The interest rate data are available from the Federal Reserve Board's H.15 release.
4. Banks' balance sheet information is available from Fitch IBCA. In particular, we retrieve

the annual information of total liabilities for each bank, and use the interpolated time series (using linear interpolation) to decide on the weight of each bank in the portfolio.

References

- Alexander, C., Sheedy, E., 2008. Developing a stress testing framework based on market risk models. *Journal of Banking and Finance* 32, 2220–2236.
- Allenspach, N., Monnin, P., 2006. International integration, common exposures and systemic risk in the banking sector: An empirical investigation. Working paper, Swiss National Bank.
- Amato, J., Luisi, M., 2006. Macro factors in the term structure of credit spreads. BIS Working Paper No. 203.
- Amato, J., Remolona, E., 2003. The credit premium puzzle. *BIS Quarterly Review* 2003-4, 51–63.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., Labys, P., 2003. Modeling and forecasting realized volatility. *Econometrica* 71, 579–625.
- Ang, A., Piazzesi, M., 2003. A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics* 50, 745–787.
- Avesani, R., Pascual, A. G., Li, J., 2006. A new risk indicator and stress testing tool: A multifactor Nth-to-default CDS basket. IMF Working Paper No. 06/105.
- Barndorff-Nielsen, O., Shephard, N., 2004. Econometric analysis of realized covariation: High frequency based covariance, regression, and correlation in financial economics. *Econometrica* 72, 885–925.
- Basurto, M. S., Padilla, P., 2006. Portfolio credit risk and macroeconomic shocks: Applications to stress testing under data-restricted environments. IMF Working Paper No. 06/283.
- Berndt, A., Douglas, R., Duffie, D., Ferguson, M., Schranz, D., 2005. Measuring default risk premia from default swap rates and EDFs. Working paper, Stanford University.
- Blanco, R., Brennan, S., March, I. W., 2005. An empirical analysis of the dynamic relationship between investment-grade bonds and credit default swaps. *Journal of Finance* 60, 2255–2281.
- Bongaerts, D., de Jong, F., Driessen, J., 2008. Liquidity and liquidity risk premia in the CDS market. Working paper, University of Amsterdam.
- Borio, C., 2003. Towards a macro-prudential framework for financial supervision and regulation? BIS Working Paper No. 128.
- Borio, C., 2006. Monetary and financial stability: Here to stay? *Journal of Banking and Finance* 30, 3407–3414.

- CGFS, 2000. Stress testing by large financial institutions: Current practice and aggregation issues. Committee on the Global Financial System Report (14), Bank for International Settlements.
- CGFS, 2005. Stress testing at major financial institutions: Survey results and practice. Committee on the Global Financial System Report (24), Bank for International Settlements.
- Chan-Lau, J. A., Gravelle, T., 2005. The End: A new indicator of financial and nonfinancial corporate sector vulnerability. IMF Working Paper No. 05/231.
- Crosbie, P., 2005. Global correlation factor structure: Modeling methodology. Moody's KMV Documents.
- Daniels, R., Koopman, S. J., Lucas, A., 2005. A non-Gaussian panel time series model for estimating and decomposing default risk. DNB Working Paper No. 55.
- Das, S., Duffie, D., Kapadia, N., Saita, L., 2007. Common failings: How corporate defaults are correlated. *Journal of Finance* 62, 93–117.
- Drehmann, M., 2008a. Macroeconomic stress testing banks: A survey of methodologies. Working paper, Bank for International Settlements.
- Drehmann, M., 2008b. Stress tests: Objectives, challenges and modelling choices. *Sveriges Riksbank Economic Review* 2/2008, 60–92.
- Driessen, J., 2005. Is default event risk priced in corporate bond? *Review of Financial Studies* 18, 165–195.
- Driessen, J., Maenhout, P., Vilkov, G., 2006. Option-implied correlation and the price of correlation risk. Working paper, INSEAD.
- Duffie, D., 1999. Credit swap valuation. *Financial Analysts Journal* 55, 73–87.
- Duffie, D., Eckner, A., Horel, G., Saita, L., 2008. Frailty correlated default. *Journal of Finance*, forthcoming.
- Duffie, D., Saita, L., Wang, K., 2007. Multi-period corporate default prediction with stochastic covariates. *Journal of Financial Economics* 83, 635–665.
- Eom, Y. H., Helwege, J., Huang, J., 2004. Structural models of corporate bond pricing: An empirical analysis. *Review of Financial Studies* 17, 499–544.
- Forte, S., Peña, J. I., 2009. Credit spreads: An empirical analysis on the informational content of stocks, bonds, and CDS. *Journal of Banking and Finance*, forthcoming.
- Gibson, M., 2004. Understanding the risk of synthetic CDOs. Working paper, Board of Governors of the Federal Reserve System.

- Goodhart, C., 2006. A framework for assessing financial stability? *Journal of Banking and Finance* 30, 3415–3422.
- Goodhart, C., Sunirand, P., Tsomocos, D., 2005. A risk assessment model for banks. *Annals of Finance* 1, 197–224.
- Goodhart, C., Sunirand, P., Tsomocos, D., 2006. A model to analyse financial fragility. *Economic Theory* 27, 107–142.
- Hancock, D., Passmore, W., 2008. Financial stability and the capital adequacy of large U.S. banking organizations: A VAR in VaR approach. Working paper, Board of Governors of the Federal Reserve System.
- Huang, J., Huang, M., 2003. How much of the corporate-treasury yield spread is due to credit risk? Working paper, Penn State University.
- Hull, J., White, A., 2004. Valuation of a CDO and an n-th to default CDS without Monte Carlo simulation. *Journal of Derivatives* 12, 8–23.
- Inui, K., Kijima, M., 2005. On the significance of expected shortfall as a coherent risk measure. *Journal of Banking and Finance* 29, 853–864.
- Jarrow, R., 2001. Default parameter estimation using market prices. *Financial Analysts Journal* 57, 75–92.
- Lehar, A., 2005. Measuring systemic risk: A risk management approach. *Journal of Banking and Finance* 29, 2577–2603.
- Longstaff, F., Mithal, S., Neis, E., 2005. Corporate yield spreads: Default risk or liquidity? New evidence from the credit-default-swap market. *Journal of Finance* 60, 2213–2253.
- Männasoo, K., Mayes, D. G., 2009. Explaining bank distress in Eastern European transition economies. *Journal of Banking and Finance* 33, 244–253.
- Merton, R., 1974. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance* 29, 449–470.
- Norden, L., Wagner, W., 2008. Credit derivatives and loan pricing. *Journal of Banking and Finance* 32, 2560–2569.
- Tang, D. Y., Yan, H., 2006. Liquidity, liquidity spillover, and credit default swap spreads. Working paper, Kennesaw State University.
- Tarashev, N., Zhu, H., 2008a. The pricing of portfolio credit risk: Evidence from the credit derivatives market. *Journal of Fixed Income* 18, 5–24.

- Tarashev, N., Zhu, H., 2008b. Specification and calibration errors in measures of portfolio credit risk: The case of the ASRF model. *International Journal of Central Banking* 4, 129–174.
- Yamai, Y., Yoshida, T., 2005. Value-at-risk versus expected shortfall: A practical perspective. *Journal of Banking and Finance* 29, 997–1015.
- Zhu, H., 2006. An empirical comparison of credit spreads between the bond market and the credit default swap market. *Journal of Financial Services Research* 29, 211–235.

Table and Figure Legends

Table 1: This table describes the summary statistics of credit factor variables and financial market variables.

Table 2: The dependent variable is the observed asset return correlation in the next quarter (between period t and $t + 12$), with the Fisher transformation applied. Explanatory variables include the current period one-quarter and one-week asset return correlations and financial market variables. The reported t-statistics (in the parenthesis) are based on Newey-West HAC covariance matrix with the truncation lag of 20. ** and * represent significance of coefficients at the 95% and 90% confidence levels respectively.

Table 3: The dependent variables are indicators of systemic risk in the banking group, including our measure of the insurance premium against distressed losses and nth-to-default probability measures with $n = 1$, $n = 2$ and $n \geq 1$. Explanatory variables are average PDs and (forward-looking) correlations $\bar{\rho}$ that are used to calculate these indicators. t-statistics are in the parenthesis and ** represents significance at the 95% confidence level.

Table 4: The results are based on a VAR analysis in which the number of lags, which equals one, is determined by the Schwarz Information Criteria. Endogenous variables include average PDs (Logit transformation applied), average one-week correlations (Fisher transformation applied), fed fund rates, term spreads, one-month returns and implied volatility of the S&P500 index. t-statistics are in the brackets. ** and * represent significance of coefficients at the 95% and 90% confidence levels respectively.

Table 5: The results show the impacts of explanatory variables, including credit risk factors and financial market variables, on PDs of individual banks (Logit transformation applied). ** and * represent significance of coefficients at the 95% and 90% confidence levels respectively.

Figure 1: This figure plots the time series of (weighted-average) risk-neutral PDs, average 1-week realized correlations and average 1-quarter future correlations. The solid lines refer to the observed data. The dash-dotted lines refer to in-sample predictions based on: (1) a VAR analysis that consists of credit risk factors (average PDs and 1-week realized correlations) and financial market variables (fed fund rates, term spreads, S&P500 one-month returns and implied volatility), as shown in Table 4; (2) regressions of individual PDs on the lagged own variable and the current-period market variables (Table 5); and (3) a regression of future (one-quarter) correlations on the current-period 1-quarter correlations, weekly correlations and other market factors (Table 2, Regression 3).

Figure 2: This figure plots the time series of fed fund rates, term spreads, S&P500 one-month returns and implied volatility (VIX). The solid lines refer to the observed data. The dash-

dotted lines refer to in-sample predictions as described in Figure 1.

Figure 3: The indicator refers to the price of insurance against banking distresses, i.e. the risk-neutral expectation of credit losses that equal or exceed 15% of the corresponding banking sector's liabilities. The prices are shown as the cost per unit of exposure to these liabilities in the upper panel and are shown in dollar term in the lower panel. The indicators are calculated based on individual risk-neutral PDs implied from CDS spreads and forecasted one-quarter asset return correlations.

Figure 4: This graph plots the movements in the indicator of systemic risk, defined as the price of insurance against banking distresses (credit losses exceed 15% of total liabilities of the banking sector), under stress testing scenarios. The stress testing adopts hypothetical scenarios based on the statistical properties of error terms in the model (described in Figure 1). In particular, the bootstrapping technique is used to simulate the future path of shock terms. The simulation is repeated 1,000 times and in each time the corresponding indicator of systemic risk is calculated. The lines plot the mean and 95% percentile distributions of the indicator.

Figure 5: This graph plots the movements in the indicator of systemic risk, defined as the price of insurance against banking distresses (credit losses exceed 15% of total liabilities of the banking sector), under stress testing scenarios. The stress testing adopts historical scenarios, i.e. using shocks in macro-financial factors as observed during the LTCM crisis (dotted line) and the September 11 episode (dashed line). The model framework is the same as the one specified in Figure 1 except that the VAR model only includes the four financial market variables.

Figure 6: The bootstrapping stress testing technique, described in Figure 4, is adopted to forecast the movements of the indicator of systemic risk (defined as in Figure 4) 12 weeks into the future. The dashed line plots the predicted average of the indicator and the dash-dotted lines the 95% confidence interval of the predicted values. The solid line refers to the *ex post* realization of these measures.

Table 1: **Summary statistics**

<i>Variables</i>	<i>Mean</i>	<i>Std. Dev.</i>	<i>Min</i>	<i>Max</i>
CDS spread (bps): bank 1	25.34	20.45	9.25	112.50
CDS spread (bps): bank 2	30.40	22.62	7.66	149.45
CDS spread (bps): bank 3	58.18	59.21	18.03	723.61
CDS spread (bps): bank 4	34.06	33.06	6.80	225.02
CDS spread (bps): bank 5	46.48	29.69	17.79	230.52
CDS spread (bps): bank 6	39.57	24.16	10.87	172.61
CDS spread (bps): bank 7	62.92	60.49	17.70	438.00
CDS spread (bps): bank 8	56.59	54.03	14.63	334.07
CDS spread (bps): bank 9	51.00	42.79	17.97	330.62
CDS spread (bps): bank 10	35.18	25.99	13.58	158.54
CDS spread (bps): bank 11	39.82	48.08	9.07	327.04
CDS spread (bps): bank 12	28.45	22.30	5.93	151.72
CDS spread (bps): weighted average	44.16	34.42	12.72	259.50
1-Week realized correlation	0.51	0.13	0.12	0.82
1-Quarter realized correlation	0.53	0.09	0.27	0.71
Fed fund rate (%)	4.86	2.13	0.96	9.90
Term spread (%)	1.58	1.17	-0.82	3.84
1-Month SP500 return (%)	0.66	4.18	-26.47	13.97
SP500 implied volatility (%)	20.39	7.65	9.04	98.81

Table 2: Forecasting asset return correlations

	Regression 1	Regression 2	Regression 3
$\tilde{\rho}_{t-12,t}$	0.52** (5.4)	0.63** (6.1)	0.52** (4.7)
$\tilde{\rho}_{t-1,t}$	0.18** (4.7)		0.12** (3.8)
FFR _t		-0.030 (-1.2)	-0.026 (-1.1)
TERM _t		-0.038 (-1.2)	-0.033 (-1.1)
SP500 ret _t		-0.0046** (-3.6)	-0.0036** (-2.9)
VIX _t		0.0015 (0.9)	0.0012 (0.8)
constant	0.19** (3.6)	0.36** (2.5)	0.33** (2.3)
Adjusted R^2	0.54	0.55	0.56
Observations	415	415	415

Table 3: Impacts of PDs and correlations on indicators of systemic risk

	Indicators of systemic risk			
	Price of insurance	$n = 1$	$n = 2$	$n \geq 1$
PD_t	0.2077** (84.0)	1.0994** (87.9)	0.3085** (159.7)	1.6952** (157.5)
$\bar{\rho}_t$	0.0029** (12.7)	-0.0204** (-17.8)	0.0008** (4.4)	-0.0157** (-15.9)
constant	-0.0021** (-17.6)	0.0145** (24.5)	-0.0005** (-5.9)	0.0110** (21.7)
Adjusted R^2	0.97	0.96	0.99	0.99
Observations	387	387	387	387

Table 4: VAR analysis

	$\tilde{P}D$	$\tilde{\rho}_W$	FFR	Term	SP500 ret	VIX
$\tilde{P}D(-1)$	0.98** (66.8)	0.055** (2.8)	-0.037* (-1.8)	0.033 (1.5)	-0.34 (-0.8)	0.66* (1.8)
$\tilde{\rho}_W(-1)$	0.083** (2.4)	0.49** (10.7)	-0.031 (-0.6)	0.026 (0.5)	0.11 (0.1)	-0.22 (-0.3)
FFR(-1)	0.010 (0.9)	-0.054** (-3.9)	0.94** (64.4)	-0.012 (-0.8)	-0.38 (-1.2)	0.084 (0.3)
Term(-1)	0.012 (0.8)	-0.071** (-3.9)	-0.064** (-3.4)	0.97** (47.8)	-0.47 (-1.1)	0.097 (0.3)
SP500 ret(-1)	-0.0025** (-2.0)	-0.0029* (-1.7)	-0.00063 (-0.4)	-0.00047 (-0.2)	0.73** (18.6)	0.0048 (0.1)
VIX(-1)	-0.00084 (-0.8)	0.0012 (0.9)	-0.0011 (-0.8)	0.0024 (1.6)	0.030 (1.0)	0.92** (35.5)
Constant	-0.18 (-1.5)	0.85** (5.4)	0.14 (0.8)	0.20 (1.2)	-0.44 (-0.1)	4.70 (1.5)
Adjusted R^2	0.97	0.43	0.99	0.99	0.53	0.91
Observations	386	386	386	386	386	386

Table 5: Determinants of individual PDs

Factors	Bank 1	Bank 2	Bank 3	Bank 4	Bank 5	Bank 6	Bank 7	Bank 8	Bank 9	Bank 10	Bank 11	Bank 12
$\tilde{P}D_{i,t-1}$	0.70**	0.63**	0.68**	0.51**	0.38**	0.71**	0.45**	0.57**	0.38**	0.81**	0.79**	0.68**
$\tilde{P}D$	0.25**	0.39**	0.36**	0.63**	0.50**	0.23**	0.63**	0.50**	0.61**	0.10**	0.29**	0.35**
$\tilde{\rho}_W$	-0.04	-0.004	0.15**	0.01	0.11**	0.13**	0.10**	0.15**	0.17**	0.02	0.03	0.05
FFR	-0.02	0.03**	0.10**	-0.03**	0.003	-0.03**	0.08**	-0.02	-0.03**	-0.0003	0.02**	0.0000
TERM	-0.02	0.04	0.08	-0.04	-0.01	-0.02*	0.05	-0.03	-0.06	0.01	0.01	0.02
SP500 ret	0.0004	-0.005**	-0.006**	-0.006**	0.001	-0.005**	-0.003**	-0.004**	-0.001	0.002	-0.003**	-0.004**
VIX	0.0002	-0.003**	-0.004**	-0.004**	0.002**	0.001	-0.004**	-0.004**	-0.003**	0.004**	-0.003**	-0.004**
constant	-0.27	-0.09	-0.17	0.78**	-0.64**	-0.31**	0.27**	0.51**	0.20*	-0.57**	0.33**	0.006
Adj- R^2	0.92	0.98	0.98	0.98	0.97	0.97	0.99	0.98	0.97	0.91	0.98	0.97
Obs.	274	381	386	386	386	385	386	386	386	363	386	386

Figure 1: Portfolio credit risk and macro-financial factors




Figure 2: Financial market factors

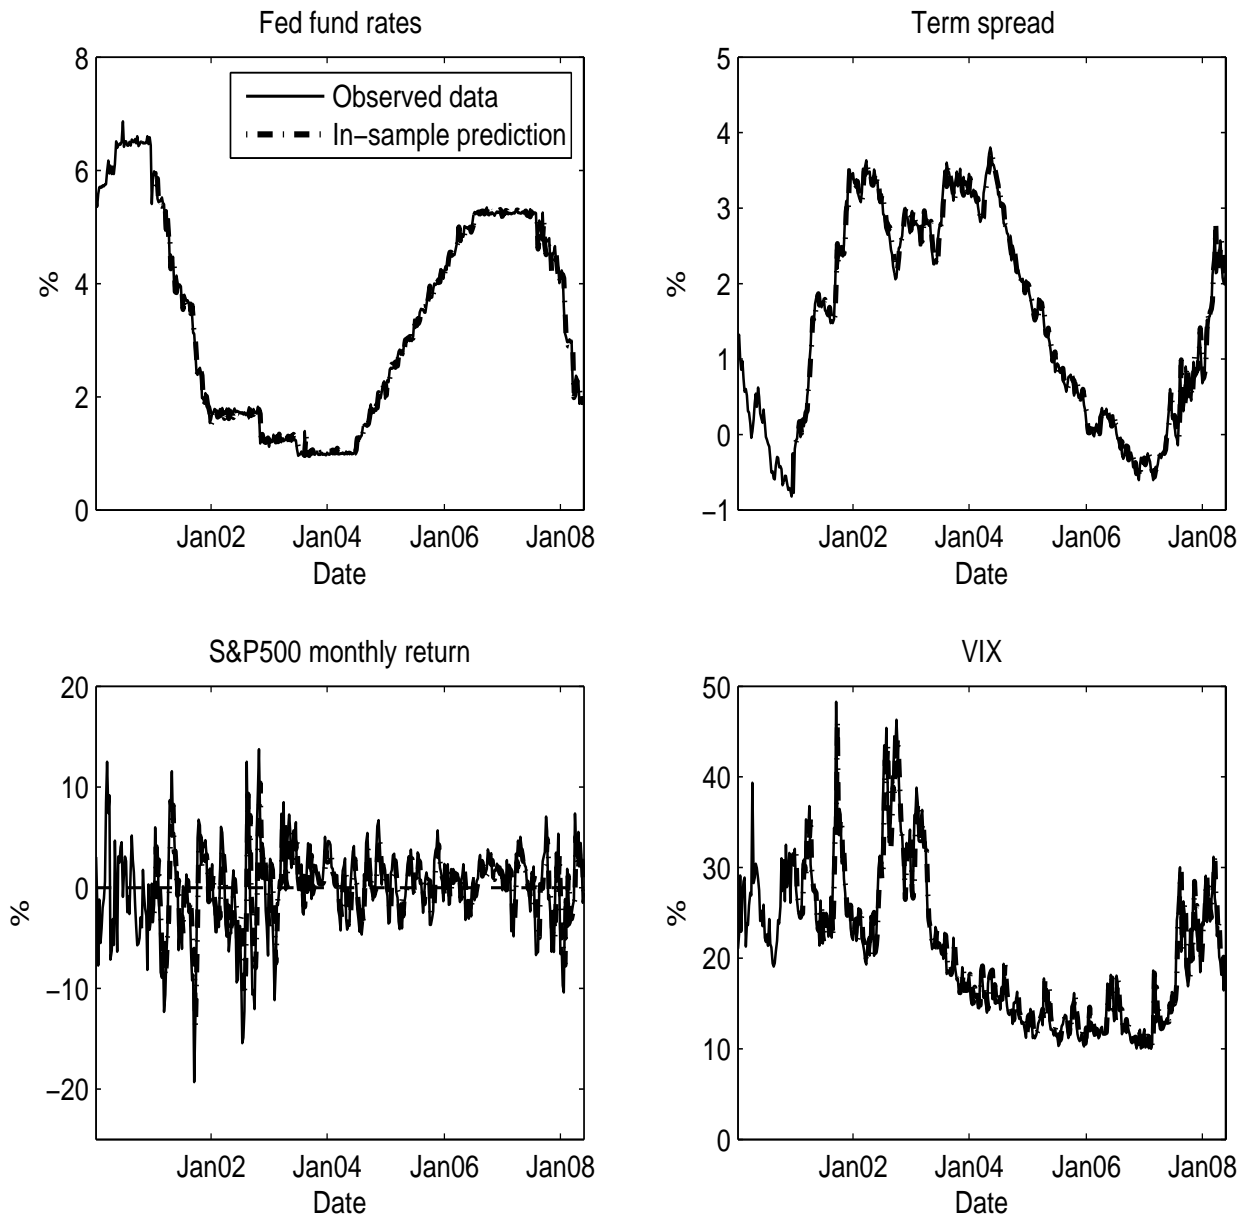


Figure 3: Indicator of systemic risk in the banking sector

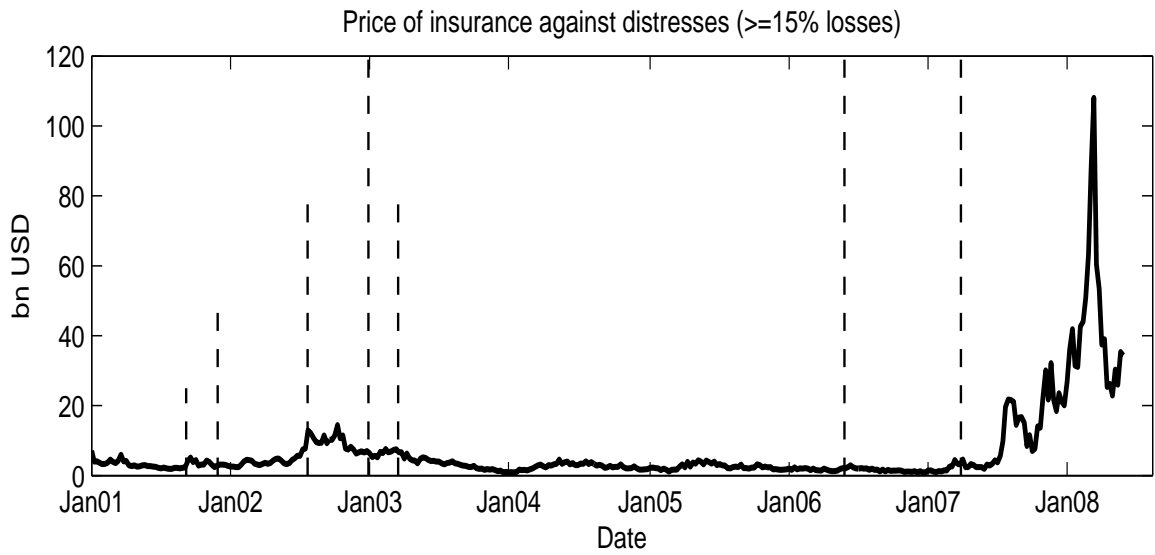
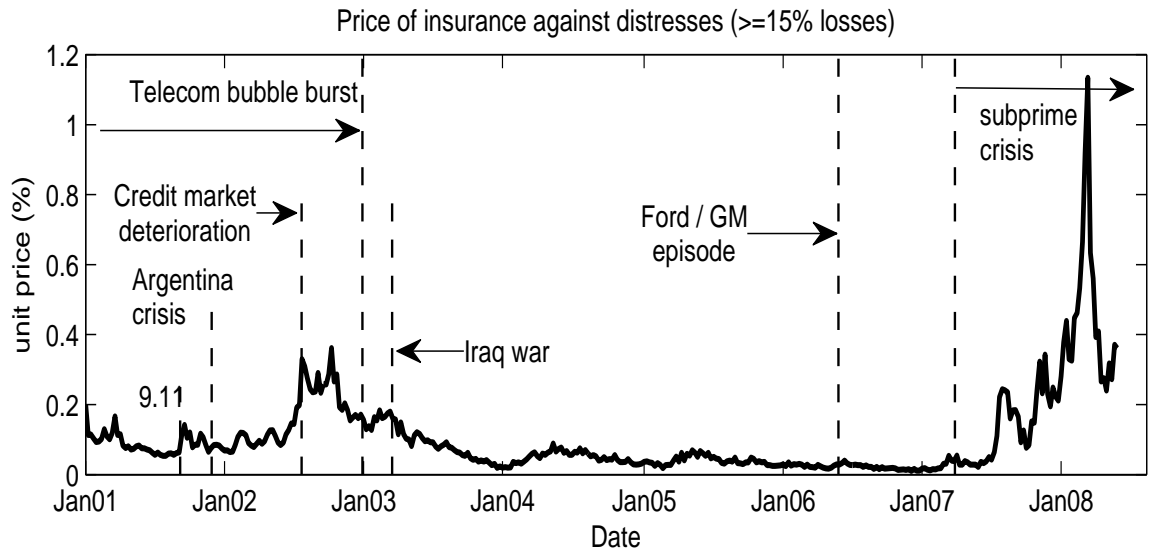


Figure 4: Stress testing exercises based on statistical scenarios




Figure 5: Stress testing exercises based on historical scenarios




Figure 6: Forecasting the systemic risk indicator

