Does Basel II Pillar 3 Risk Exposure Data help to Identify Risky Banks?

Ralf Sabiwalsky*

* Freie Universität Berlin, Germany

ECONOMIC RISK   BERLIN

SFB 649

Does Basel II Pillar 3 Risk Exposure Data help to Identify Risky Banks?

Ralf Sabiwalsky[*]

February 2, 2012

Abstract

Basel II Pillar 3 reports provide information about banks' exposure towards a number of risk factors, such as corporate credit risk and interest rate risk. Previous studies find that the quality of such information is likely to be weak. We analyze the marginal contribution of pillar 3 exposure data to the quality of equity volatility forecasts for individual banks. Our method uses (local in time) measures of risk factor risk using a multivariate stochastic volatility model for five risk factors, and uses measures of bank sensitivity with respect to these risk factors. We use two sets of sensitivity measures. One takes into account pillar 3 information, and the other one does not. Generally, we generate volatility forecasts as if no market prices of equity were available for the bank the forecast is made for. We do this for banks for which such data is, in fact, available so that we can conduct ex post - tests of the quality of volatility forecasts. We find that (1) pillar 3 information allows for a better-than-random ranking of banks according to their risk, but (2) pillar 3 exposure data does not help reduce volatility forecast error magnitude.

Keywords: Risk Reporting, Stochastic Volatility, Risk Factors
JEL Classification: G17, G21

1 Introduction

Banks' stock returns are driven by systematic risk and by idiosyncratic risk. The proportion of systematic risk is likely to be substantial since banks' assets are diversified and sensitive with respect to systematic risk factors. Among such factors are corporate equity returns, interest rates and aggregate credit risk. Thus, knowing the risk inherent in such risk factors should make it possible to make good forecasts of system-wide bank stock return volatility. However, there is also individual variation in bank volatility, too. Determinants of this individual variation include bank-specific sensitivity with respect to systematic risk factors on the one hand and idiosyncratic risk on the other hand. In Europe, banks are obliged by European law to publish standardized reports on their risk exposure (Basel II pillar 3 reports). Because of this, we analyze a sample of European banks. This article aims at finding out whether pillar 3 disclosures have the potential to improve estimates of bank-specific future volatility.

Estimates of the future volatility of a bank's equity returns are useful for a number of problems. These problems include (a) investment decisions with regard to the bank's equity, (b) interbank lending decisions, where equity return risk is an indicator of default risk due to the strong impact of asset quality and risk management practices on the return distribution, (c) regulators' assessment of the stability of the financial system and (d) the valuation of derivatives written on the bank's equity. For a number of banks, empirical volatility and option-implied volatility can be used as estimates of equity volatility. However, for parts of banks, in adverse market conditions (e.g. due to low liquidity or trading restrictions) or for validation purposes, an analyst might wish or have to estimate volatility using information which does not rely on market data relating to the bank in question. Furthermore, there are a number of banks for which such data is not available at all. In such situations, using knowledge on the risk inherent in systematic risk factors and about the exposure of a bank with respect to these risk factors could be useful for assessing bank-specific risk. Whether this is indeed the case is an empirical question, and this article is an attempt answer it.

On the one hand, pillar 3 disclosures on risk exposure should help differentiating between banks that are very sensitive with respect to certain risk factors and banks which are robust against risk factor volatility. Consider two banks, one of which hedges most of its net interest rate risk (low interest rate market value at risk), while the second one leaves open most of its net interest rate risk (high interest rate market value at risk). When interest rate volatility increases, the first bank would exhibit only a modest increase in equity volatility, while the second bank would exhibit a substantial increase. On the other hand, while pillar 3 rules ensure a minimum level of comparability between disclosures, the methodologies that different banks use for calculating risk disclosures vary. Consider the situation where differences between measures of risk exposure are larger across different calculation methodologies than across banks. Then using pillar 3 disclosures could make volatility forecasts using individual risk exposure data worse than forecasts which ignore such data. Our results indicate that the latter situation is true.

2 Priors on the Usefulness of Banks' Disclosures about Risk Exposure

Operating a bank means taking on financial risks. Typically, the methodologies developed for managing financial risks and the internal organization of banks' risk management units reflect a distinction between market risk and credit risk which is stricter than the interrelation between these two risk types would suggest. This has resulted in the use of risk measures which separately assess market risk and credit risk. For the purposes of this study, we define risk as the extent to which the future value of a bank's equity is stochastic. We define financial risk (as opposed to strategic, operational, legal and other risks) as the stochasticity of future equity returns which results from sensitivity of the bank's future cashflows towards changes in the values of financial contracts that the bank has already entered into. We further distinguish between sources of financial risks that go beyond the distinction between credit risk and market risk. We focus on large banks operating in the retail, corporate and wholesale markets. We concentrate on five sources for the risk of changes in the value of a bank's financial contracts:

1. changes of the likelihoods that corporate borrowers will be able to fulfil their obligations (corporate credit risk),

2. changes of the likelihoods that sovereign borrowers will be able to fulfil their obligations (sovereign credit risk),

3. changes of interest rates (interest rate risk),

4. changes of stock prices of nonfinancial companies (stock market risk),

5. changes to the value of mortgage backed securities (MBS market risk).

Assessing the total risk of a bank can be accomplished in two or three steps. The first step will consist of (a) an assessment of the vulnerability of the bank's existing contracts with respect to these five (and possibly other) risk factors and (b) an assessment of the capability of the bank's risk management methodology to handle the risks arising from contracts that the bank will enter into in the future. The second step will require an analysis of idiosyncratic risks which arise, in particular, from the bank's competitive position. In addition, possibly as a third step, an analysis of the regulatory environment will be helpful, and regulation-induced risks to a bank's future profitability will hit all banks in the same jurisdiction, but different banks may be differently well prepared to cope with regulatory changes.

Banks are obliged to publish financial reports and make regulatory disclosures. The aim of such publications is, among others, to inform creditors about their financial risk position. That is, such publications should help financial analysts and rating agencies to find out how vulnerable the bank is with respect to changes in risk drivers. We argue that the five risk factors listed above represent very important risk drivers. Now consider the goal to estimate the volatility of a bank's equity returns over one period ahead, where this could be a day, a month, a year or some other time interval of interest. We consider this to be an elementary task an analyst wishes to complete as part of a bank risk analysis.

We ask the question whether disclosures made by banks, following state of the art methodology, help the analyst in completing this task. We restrict our analysis to the data that can be used to assess the sensitivity of the bank with respect to the five risk drivers listed above, i.e. we explicitly exclude the analysis of idiosyncratic risk. The usefulness of risk disclosures is a critical information for analysts who use such information to take investment decisions. Furthermore, regulators' decisions concerning regulatory measures require estimates of their potential effect on banks' risk taking behaviour. Banks' disclosures about risk should only be used for such an assessment when they are of sufficient quality. That is, when potential inconsistencies between the risk measurement and disclosure methodologies among different banks and / or over time are smaller in magnitude than the actual differences in risk (see González, 2005 on the analysis of the relationship between regulatory measures and banks' risk-taking behaviour).

In the remainder of this section, we will first discuss why this is an important question to ask, i.e. we present arguments in support and arguments in contradiction to the hypothesis that mandatory risk disclosures are useful for volatility estimation. Second, we will present a methodology which allows for an assessment of the usefulness of risk disclosures.

To facilitate the discussion, we will first specify the risk disclosure data that we consider helpful for an assessment of banks' financial risks. First consider loans given out to borrowers of varying credit quality. A deterioriation in average probabilities of default over all rating classes will lead to a severe deterioration of the expected net present value of interest and principal payment from borrowers in low quality rating classes, while for high rating classes, the effect will be much less pronounced. Thus, for assessing the sensitivity of the bank's equity with respect to changes in credit quality, we suggest that the credit exposure of a bank in non-investment grade rating classes is a valid indicator. In their Basel II pillar 3 report, banks have to disclose by rating class their credit exposure. The disclosed exposures must be broken down by borrower type, which implies that the non-investment grade corporate and sovereign credit exposure are disclosed separately (see BIS 2007, part 4, sec. II, Table 5b and 6d). We hypothesize that

H1 the higher a bank's corporate credit exposure in non-investment grade rating classes, the more vulnerable its equity returns are with respect to aggregate corporate credit risk and

H2 the higher a bank's sovereign credit exposure in non-investment grade rating classes, the more vulnerable its equity returns are with respect to aggregate sovereign credit risk.

We now turn to interest rate risk, which generally arises from trading activities and from banking book transactions. Many large banks transfer banking book interest rate exposure to the trading book by internal transactions, such that total interest rate risk is solely managed through managing the net interest rate exposure in the trading book. For such banks, and only for them, interest rate risk in the banking book can be read from risk disclosures about interest rate risk in the trading book. Banks can choose to hedge most of their interest rate exposure or to leave open the net position, depending on the market risk management strategy. When a bank's strategy is to hedge most of its interest rate exposure, interest rate volatility will not translate into substantial variation in the value of its financial contracts for this bank. Now consider the components of the value of a bank's equity.

4

Part of it results from financial contracts the bank has already entered into. When interest rate risk is fully hedged, this component is not sensitive with respect to interest rate volatility. Part of the equity value, however, results from the expectation that the bank will be able to enter into profitable contracts in the future. The interest rate risk arising from future contracts will not generally be fully hedged, even if the strategy is to hedge interest rate risk from existing contracts. Banks use different measures for interest rate sensitivity. They present the change in value of existing contracts that would arise from prespecified changes of certain interest rates (sensitivity measures). Further, they disclose a quantile of the profit/loss distribution of their trading book under the assumption that variability is only induced by interest-rate risk (VaR measures), as required by BIS (2007), part 4, sec. II, Table 11e.[1] Sensitivity measures cannot be used to compare different banks because the interest rate changes that they use are not generally identical. In contrast, VaR disclosures could by comparable, because pillar 3 regulations require the use of a set of predefined parameters for the analysis. However, for some banks, banking book interest rate is not made part of the VaR analysis. Disclosures for interest rate risk which remains in the banking book follow the management approach and thus are not useful for comparing more than a very small number of banks, since costly imprecise adjustments are required for making such disclosures comparable (BIS 2007, part 4, sec. II, Table 14b requires to publish income shocks that would result from interest rate shocks of magnitudes identical to those being used for management's analysis of this risk). When we restrict our focus on the component of equity value that arises from a positive net present value of the set of existing contracts, we can hypothesize that

H3 the higher a bank's exposure to interest rate risk, measured by the interest rate VaR, the more sensitive its equity value is with respect to changes of interest rates.

A similar conclusion can be drawn for equity market risk. We suggest that changes in the level of aggregate stock indices leads to changes of the equity value of a bank in as far as the bank has open trading positions subject to stock market risk. The extent to which a bank has such open positions could possibly be inferred from observing its equity market risk VaR from trading activities, which is disclosed according to BIS (2007), part 4, sec. II, Table 11e. Further insight could be gained by observing the exposure to equity in the banking book (BIS 2007, part 4, sec. II, Table 13e). We can (and do) exclude stock market risk specific to the financial sector when defining the time series for risk factor no. 4 (stock market risk), because otherwise we would measure volatility which is partly driven by the idiosyncratic risk of the banks in our sample. We cannot exclude stock market risk specific to the financial sector when assessing the equity market VaR for a bank, since only aggregate stock market VaR are disclosed. We hypothesize that this inavoidable imprecision is not material, so that we can write down

H4 the higher a bank's exposure to stock market risk, measured by the stock market VaR, the more sensitive its equity value is with respect to aggregate nonfinancial stock market returns.

Mortgage backed securities have attracted considerable attention following the financial crisis of 2007/2008. This was due to substantial holdings of such instruments in bank's books. Many banks provide detailed

---

[1]Generally, pillar 3 requirements are binding law in the European Union, see European Parliament and European Council (2006), where "binding" means that member states are obliged to introduce them into national law.

disclosures about their exposure to such instruments. Yet, while the data presented is detailed, it is difficult to obtain a measure of bank's exposure to MBS that is comparable across all banks without a cost-intensive analysis of each individual bank's report. Pillar 3 requires the disclosure of exposure in securitizations by rating class. We argue that the exposure to non-investment grade securitization positions provides an indicator for the vulnerability of banks' equity value with respect to MBS market volatility, since a substantial portion of securitizations exposures pertain to MBS transactions. Securitization disclosures follow BIS (2007), part 4, sec. II, Table 9g. We hypothesize that

H5 the higher a bank's exposure to non-investment grade securitizations, the more sensitive it is with respect to volatility in the MBS market.

The aggregate hypothesis we state is

H6 volatility forecasts that take into account bank - specific risk exposure data are better than volatility forecasts that do not.

At first sight, the usefulness of such risk disclosures for volatility estimation - i.e. the validity of hypotheses H1 to H5 - should be self-evident. A more detailed analysis, however, reveals that there is substantial prior evidence against their validity. We will first discuss the quality of disclosures regarding exposure to low-quality credit risk. Then, we will discuss the quality of VaR disclosures in their role as an indicator of (a) the risk factor - specific sensitivity of the value of a banks' portfolio of existing contracts and of (b) a bank's strategy to manage factor-specific risk, i.e. whether the bank has the general goal to hedge net exposure to a risk factor or not.

Most large banks operating under the Basel II regime apply internal ratings systems to assign exposure to rating classes. Internal ratings systems require approval by supervisory bodies before their output can be used to calculate risk-weighted assets. This ensures a minimum level of coherence between systems used by different banks. Yet, it is far from certain that the same borrower will be assigned the same rating by two different banks. This is because first, the size and quality of historical data used for calibrating ratings systems varies among banks. Second, given a historical dataset, rating analysts enjoy methodological discretion when designing rating models. Third, before a rating is calculated for a specific borrower, the bank's analyst applies adjustments to the borrower's financial data, and these adjustments follow both bank-specific general rules and also considerations specific to the individual analyst. Fourth, many ratings systems use qualitative data to account for characteristics of the borrower that are not measurable by financial data but at the same time exhibit predictive power for defaults. As an example, the analyst might be asked to assign a grade to the ability of the management of the borrower. Such assessment is subjective and unlikely to result in coherent results among different lenders. As a consequence, it remains an empirical question whether disclosures on rating-specific exposure data by two banks is coherent enough to allow for a comparison of these two banks' credit risk sensitivity. Cary and Hrycay (2001) analyze the quality that can reasonably be expected for rating systems. They argue that ratings are subject to bias, instability and gaming (see also Jacobson et al., 2006). In line with their arguments, empirical evidence suggests that the potential for meaningful comparisons between banks is limited. Carey (2001) compares ratings of

firms that have been assigned a rating by several U.S. banks. He finds variation in assigned ratings. Jacobson et al. (2006) analyze the rating systems of two Swedish banks and firms who have borrowed from both banks and find that internal rating systems at different banks provide different ratings in a substantial number of cases.

Value at Risk estimates require reasonable knowledge about the shape of the profit/loss distribution far away from the center of the distribution. Banks can choose from different methodologies to extract knowledge on this distribution from data on existing contracts and from historical and from current market pricing data. A large number of different approaches to modelling relevant stochastic processes and stochastic dependencies exist. Supervisors ensure a minimum quality of internal VaR models by requiring an extra capital charge if backtesting results are weak, but they do not require that for a unique portfolio of transactions, models from different banks provide the same VaR result. Furthermore, banks do not necessarily have incentives to report accurate VaR. On the one hand, Lucas (2001) argues that due to Basel II regulations, it can be rational to understate VaR for some banks. On the other hand, underestimated VaR result in a high number of exceptions to be reported to the supervisor and possibly to the capital market. Pérignon et al. (2008) argue that thus, banks might overestimate VaR as this preserves high reputation. As a consequence, it is not unlikely that two banks would obtain two substantially different VaR estimates for the same portfolio of contracts. This assertion is reinforced by the finding that the data used for backtesting VaR models is not generally appropriate, see Frésard et al. (2011). Furthermore, banks continuously strive to improve their methodology for measuring risk. Thus, different VaR estimates might be calculated for one bank at two different points in time for a fixed hypothetical portfolio of contracts. If this were the case, it would be detrimental to the usefulness of risk-factor specific VaR estimates in the assessment of a bank's sensitivity with respect to the risk factor. Previous empirical evidence, again, is not very promising. Pérignon et al. (2008) consider VaR disclosures of several banks and argue that these overstate the true VaR by between 19% and 79%. Since the magnitude of overstatement need not be equal among different banks, VaR of different banks might not be comparable. They further argue that there is considerable evidence that VaR models fail to accurately predict quantiles of the loss distribution (see Cassidy/Gizycki 1997, Berkowitz/O'Brien 2002 and Berkowitz et al. 2006). Yet, overstated VaR do not imply that monitoring of trading risks is generally conservative. Wong (2008) shows that even though VaR might be overstated, tail risk can be severe. In particular, he finds that four out of six large commercial US banks exhibit trading loss distributions with "fatter-than-normal" tails.

Pérignon/Smith (2010a) find that VaR disclosure standards have improved over time between 1996 and 2005. However, accuracy of VaR disclosures, in terms of the fit between implied frequency of exceedances and observed frequency of exceedances has not improved over time. Furthermore, Pérignon/Smith (2010b) compare the magnitude of the diversification effect as disclosed by US banks to the magnitude of the diversification effect as empirically estimated in their paper. They find that the disclosed values and their empirical estimates match. Since they have, however, in the companion paper (2010a) showed that disclosed VaR are higher than what empirically observed trading profit/loss distributions imply, they conclude that the disclosed VaR are biased.

From this analysis we conclude that it is an empirical question whether the quality of credit risk exposure and VaR disclosures is high enough to render this information useful for the task of forecasting individual bank equity volatility or to compare two banks in terms of their risk sensitivity. If there were only one risk factor, then banks with higher exposure should be more risky than banks with lower exposure, and the analysis would be straightforward. There are, however, a number of relevant risk factors. When bank A has higher exposure with respect to risk factor 1, but bank B has higher exposure with respect to risk factor 2, then the analysis must make use of data on the relative risk associated to both risk factors. An obvious attempt to account for that would be to follow the procedure applied in value relevance research, which aims at finding out how useful the disclosure of certain components of accounting income is. This is typically accomplished by testing ex post for correlation between accounting income ratios (or unexpected income surprises) and stock returns (see Holthausen/Watts, 2001 and Barth et al., 2001). However, we are interested in the question whether risk disclosures are *materially useful* for the purpose its target group is likely use it for. Therefore, we will apply the following methodology. The main idea is: disclosed risk measures are useful if they help an analyst to assess the risk for changes of the value of assets and/or liabilities which eventually translate into net changes of the value of its equity. Provided the bank is listed on the stock exchange, a simple measure for this risk is volatility of the bank's stock returns over future periods. However, if the bank is listed on the stock exchange, a detailed analysis of risk disclosures might not be so relevant since market data (historical returns, option-implied volatilities) can provide reasonable estimates of future volatility. When, in contrast, a bank is not listed on the stock exchange, then its risk disclosures are a major source of information about the bank's risk. An analyst who wishes to assess a nonlisted bank's future equity volatility could use risk disclosure data to find out about the bank's sensitivity with respect to important risk factors and about the bank's idiosyncratic risk. Then she could combine this knowledge with good estimates of the future volatility of the risk factors to obtain an estimate of the bank's future equity volatility *even if it is not listed on the stock exchange*. If this volatility is high, there is a high likelihood that the bank's equity will drop below regulatory minimum capitalization requirements, which could eventually trigger a default. This is exactly what we do: We use a sophisticated approach to rank banks that are listed on the stock exchange according to their future estimated equity volatility, taking into account risk disclosures, *but without using capital market data about these banks*. That is, we treat these banks as if they were not listed on the stock exchange. We call these banks the forecast sample, and we compare our forecasts to the realized volatility for these banks. Then we use a related approach to estimate future equity volatility, but here we ignore *all information* that is specific to the banks in the forecast sample. If using pillar 3 disclosure data results in better volatility estimates compared to estimates that ignore pillar 3 information, then we would conclude that pillar 3 information provides useful information. We proceed as follows: First, we collect data on the exposure of 12 large European stock-listed banks with respect to five risk factors: corporate credit spreads, sovereign credit spreads, equity market (excluding financials) risk in the trading book, interest rates and the U.S. mortgage backed securities market. Second, we estimate the parameters of a model for the joint dynamics of the risk factors in a multivariate stochastic volatility model (the "factor model"). Third, we divide the sample into 6 (learning sample) + 6 (forecast sample) banks. For the learning sample, we estimate a factor stochastic volatility model (the "stock model"), where we estimate the

8

sensitivity of bank equity returns with respect to these risk factors and with respect to the interaction between these risk factors and risk factor exposure. We estimate the complete model and a restricted variant of the model where interaction parameters are restricted to be zero. Generally, all sensitivity parameters are restricted to be identical for the 6 "learning" banks. Fourth, we use the estimated dynamics of the risk factors and the estimated sensitivity parameters to forecast equity volatility for the forecast sample banks. We assess the quality of these volatility forecasts for the complete model and the restricted model. We forecast daily volatility over monthly intervals, but not entire return distributions, since we have only about 3 years of data (pillar 3 reports have not been published before 2008, and previous (voluntary) disclosures were so heterogeneous such that a reasonable empirical assessment is likely to be unreliable).

3 Models

We entertain two models. The first one is a multivariate model for risk factor returns. The second one is a factor model for the stock returns. This two-layer structure has been chosen upon the observation that there exist a small number of factors that substantially influence bank stock returns and that these factors are not stochastically independent. For an introduction to volatility modelling in general, and stochastic volatility modelling in particular, see Andersen et al. (2006).

3.1 Factor Model

The measurement equation for the (demeaned) daily returns of F risk factors is

$$c_t = H_t R \varepsilon_t, \tag{1}$$

where c_t is an $F \times 1$ vector of logreturns and R is an $F \times F$ lower triangular matrix with ones on the diagonal and (possibly nonzero) entries $r_{ij}, 1 \leq j < i \leq F$ below the diagonal. ε_t is an $F \times 1$ vector of normals, where both its entries are independent and also $\varepsilon_t, \varepsilon_{t'}$ are independent whenever $t \neq t'$. This structure allows for straightforward estimation of the correlations between factors, since R is the Cholesky decomposition of the covariance matrix of the factor disturbances where the parameters below the diagonal are entirely free, and at the same time the implied covariance matrix $R_t R_t'$ is positive definite. H is a diagonal matrix with the latent volatility states on the diagonal, i.e.

$$H_t := \begin{pmatrix} \exp(0.5x_{t,1}) & 0 & \ldots & 0 & 0 \\ 0 & \exp(0.5x_{t,2}) & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & \exp(0.5x_{t,F-1}) & 0 \\ 0 & 0 & \cdots & 0 & \exp(0.5x_{t,F}) \end{pmatrix}.$$

$(x_{t,i})_{i=1,...,F} =: x_t$ is the F - dimensional latent log-variance state. Its evolution is specified as

$$x_t = \mu_x + \phi_x(x_{t-1} - \mu_x) + \Gamma Q \eta_t. \tag{2}$$

μ_x is an $F \times 1$ vector of long-term mean log-variances, $\phi_x \in (0,1)$ is adjustment speed towards these means, Γ is a diagonal matrix which contains the volatilities of volatility $\sigma_j > 0$, $j = 1,...,F$ on its diagonal. η_t is an $F \times 1$ vector of normals, where its entries are independent and also $\eta_t, \eta_{t'}$ are independent whenever $t \neq t'$. In general, (contemporaneous) leverage effects can be modelled by specifying the joint distribution of ε_t and η_t as (in block notation)

$$\begin{pmatrix} \varepsilon_t \\ \eta_t \end{pmatrix} \sim \mathcal{N} \left[0, \begin{pmatrix} I & L \\ L & I \end{pmatrix} \right],$$

where L is diagonal with entries $\lambda_j, j = 1,...,F$ on the diagonal, where $\forall j = 1,...,F: \ \lambda_j \in (-1,1)$. Q is the (lower triangular) Cholesky decomposition of the covariance matrix of the state disturbances. By including elements below the diagonal, correlated volatility evolutions can be specified. In the model used for the empirical analysis, however, we specify $L = Q = I.$[2]

3.2 Stock Model

Two variants of the factor stochastic volatility model for a number of S stock returns are specified; an unrestricted variant and a restricted variant. Note that risk factors are indexed by $j = 1,...,F$ and bank stocks are indexed by $i = 1,...,S$. Conditional on the risk factor returns, returns of any two stocks are assumed independent. Thus, the volatility states in the stock model are associated purely to idiosyncratic volatility, and, we can specify the univariate dynamics separately for each stock. The unrestricted measurement equation for the (demeaned) daily return of stock i over period t, $y_{t,i}$, is

$$y_{t,i} = \sum_{j=1}^{F} (\beta_j + \xi_j a_{t,i,j})c_{t,j} + exp(0.5 v_{t,i})\epsilon_{t,i}. \tag{3}$$

β_j measures the average sensitivity of stock returns with respect to risk factor j across all stocks in the learning sample. $a_{t,i,j}$ is the exposure of bank i with respect to risk factor j as reported in the pillar 3 report for the last disclosure period ending before period t. ξ_j measures the interaction between the returns of risk factor j and banks' exposure to that risk factor, as an average across all banks in the learning sample. $v_{t,i}$ is the latent log-variance state. ϵ is standard normal, independent over time and across stocks. The dynamics of the state are modelled as

$$v_{t,i} = \mu_{v,i} + \phi_v(v_{t-1,i} - \mu_{v,i}) + \sigma_{v,i}\omega_{t,i}, \tag{4}$$

where $\text{Corr}(\epsilon_{t,i}, \omega_{s,k}) = \begin{cases} \rho & \text{if } i = k \text{ and } t = s, \\ 0 & \text{in all other cases.} \end{cases}$

---

[2]It is left for later versions of this paper to find out whether relaxing this restriction would improve model performance.

$\mu_{v,i}$ is long-term mean log-variance, $\phi_{v,i} \in (0,1)$ is adjustment speed towards this mean, and $\sigma_{v,i} > 0$ is volatility of volatility. In the restricted variant, $\forall j = 1, ..., F : \xi_j = 0$.

3.3 Parameter and State Estimation

3.3.1 Overview

A particle Markov chain Monte Carlo method is applied for estimating the parameters of the models (see Andrieu et al., 2010, for an overview on the technique). Three model components (factor model, unrestricted stock model, restricted stock model, see sections 3.1 and 3.2) are estimated separately. The procedure proceeds following the idea of Pitt et al. (2010). The approach has been chosen since it allows for the establishment of a fixed estimation framework that can flexibly handle a large set of different variants of state space models. All parameters are transformed onto the real line (for ϕ, the logit transformation is used, for correlation parameters, minus 1 plus 2 times a logit variable is used, and for strictly positive parameters, their logarithm is estimated). The idea is: we implement the Metropolis Hastings algorithm (MH) to draw samples from the posterior distribution of the parameters, and for evaluating the simulated likelihood in each iteration of the MH, we use the auxiliary particle filter (ASIR) from Pitt and Shephard (1999). We use a single component MH, i.e. we update one component of the parameter vector at a time.

In the following, the estimation procedure is described for a general state space model. We use y for the observations and v for the latent states as in the stock model, but the method applies also to the factor model where c denotes the observations and x denotes the latent states.

We use $p(\cdot|\cdot)$ for a generic conditional density defined by its arguments. The measurement density is given by

$$p(y_t|v_t, \theta) \tag{5}$$

and the state transition density is given by

$$p(v_t|v_{t-1}, \theta), \tag{6}$$

where θ is the vector of D parameters (the composition of θ depends on the model component whose parameters we estimate). We further specify a prior distribution on the parameters, $p(\theta)$, and an initial distribution for the states, $p(v_0|\theta)$. Our priors are the normal distribution with mean 0 and variance 10 for $\mu_.$, $r_.$, $\beta_.$ and $\xi_.$; a truncated normal between 0 and 1 with location 0.9 and scale 0.1 for $\phi_.$ and a truncated normal between -1 and 1 with location 0 and scale 10^6 for ρ; and an inverse gamma with parameters 0.01 and 0.01 for $\sigma_.^2$. $p(v_0|\theta)$ is a uniform distribution around $\mu_.$. Parameters are estimated by drawing samples from their posterior distribution using the MH. We use the notation $y_{1:t}$ whenever we mean $\{y_1, ..., y_t\}$, and this notation also applies to other variables (such as $x_{1:t}$). We aim at drawing from the distribution $p(\theta|y_{1:T})$. Generally, we would thus define a Markov chain with stationary distribution proportional to $p(y_{1:t}|\theta)p(\theta)$. However, the likelihood $p(y_{1:t}|\theta)$ is not known

11

in closed form due to the latent states. Thus, we use a simulated likelihood which is obtained from filtering the state space model. We use the well-known auxiliary particle filter (ASIR) from Pitt and Shephard (1999).

3.3.2 Single Component Adaptive Metropolis Hastings Algorithm

We denote the vector of parameters in the g-th iteration of the Markov chain by θ_g. Throughout the MH algorithm, we separately make one-dimensional proposals for the elements in θ. We denote the $m - th$ element of θ_g by $\theta_{g,m}$, and we define $\theta_{g,-m} := (\theta_{g,1}, ..., \theta_{g,m-1}, \theta_{g,m+1}, ..., \theta_{g,D})$ and $\theta^+_{g,-m} := (\theta_{g,1}, ..., \theta_{g,m-1}, \theta_{g-1,m+1}, ..., \theta_{g-1,D})$. The Metropolis Hastings algorithm proceeds as follows for $g = 1, .., G$ where θ_0 can be chosen arbitrarily.

$\forall g = 1, ..., G; \ \forall m = 1, ..., D :$

1. draw $\tilde{\theta}_m$ from proposal density $q_g(\theta_m|\cdot)$,

2. run the ASIR and calculate the simulated likelihood,

3. define

$$
\theta_{g,m} := \begin{cases} \tilde{\theta}_m & \text{w. prob. } \alpha_{g,m} = \min\left\{1, \frac{\hat{p}(y_t|y_{1:t-1}, \tilde{\theta}_m, \theta^+_{g,-m})p(\tilde{\theta}_m)}{\hat{p}(y_t|y_{1:t-1}, \theta_{g-1,m}, \theta^+_{g,-m})p(\theta_{g-1,m})} \cdot A_{g,m}\right\}, \\ \theta_{g-1,m} & \text{w. prob. } 1 - \alpha_{g,m}. \end{cases}
$$

The proposal density $q_g(\theta_m|\cdot)$ changes over the course of the MH (hence its index) and so does $A_{g,m}$, a generic expression by which we denote the ratio of proposal densities. For small g, we choose random walk proposals, and for large g - when something has been learnt about $p(\theta|y)$ - we choose independent sampling. In particular, we adapt the scheme suggested and described in detail by Pitt et al. (2010) to the single component MH.

We start with the adaptive random walk of Roberts and Rosenthal (2009) for the first 800 iterations. Hence, $\forall g = 1, ..., 800; \ \forall m = 1, ..., D : \ A_{g,m} = 1$. We draw from

$$
q_g(\theta_m|\cdot) = q_g(\theta_m|\theta_{g-1,m}, \theta^+_{g,-m}) = \delta_g\Phi(\theta_m|\theta_{g-1,m}, \kappa_1 I, \theta^+_{g,-m}) + (1-\delta_g)\Phi(\theta_m|\theta_{g-1,m}, \kappa_2\Sigma_{2,g-2}, \theta^+_{g,-m}). \tag{7}
$$

Here, $\kappa_1 = 0.01/D$, $\kappa_2 = 2.38^2/D$ and $\Sigma_{2,g-2}$ is the sample covariance matrix of $\theta_{1:g-2}$. Now let z be an arbitrary vector of dimension D with elements $z_m, m = 1, ..., D$ and let b be an arbitrary vector of dimension $D - 1$. Then we define $\Phi(z_m|\mu, \Sigma, b)$ as the conditional normal density for z with mean vector μ and covariance matrix Σ evaluated at z_m conditional on $(z_1, ..., z_{m-1}, z_{m+1}, ..., z_D)' = b$. We set the weight $\delta_g = 1$ until 400 proposals have been accepted; then, up to the $800th$ acceptance, we choose $\delta_g = 0.05$. From the $801th$ acceptance, we use an independent sampling scheme, where

$$
q_g(\theta_m|\cdot) = q_g(\theta_m|\gamma_g, \varphi_g, \theta^+_{g,-m}) = \sum_{l=1}^{4} \gamma_{g,l}\mathrm{f}_l(\theta_m|\varphi_{g,l}, \theta^+_{g,-m}) \tag{8}
$$

and where $A_{g,m} = \frac{q_g(\theta_m|\gamma_g,\varphi_g,\theta_{g,-m}^+)}{q_g(\bar{\theta}_m|\gamma_g,\varphi_g,\theta_{g,-m}^+)}$. That is, we use a mixture of four proposal densities each denoted by f_l, $l = 1,...,4$; where the parameters of any f_l are stacked in $\varphi_{g,l}$. Until the $1200th$ acceptance, we set $\gamma_{g,1} = 0.8$, $\gamma_{g,2} = 1 - \gamma_{g,1}$, $\gamma_{g,3} = \gamma_{g,4} = 0$, $f_1(\theta_m|\varphi_{g,l}) = \Phi(\theta_m|0,\Lambda_{1,g},\theta_{g,-m}^+)$ and $f_2(\theta_m|\varphi_{g,l}) = \Phi(\theta_m|0,10\cdot\Lambda_{1,g},\theta_{g,-m}^+)$. $\Lambda_{1,g}$ is the sample covariance matrix of the sampled parameter vectors $\theta_1,\theta_2,...$ updated at regular intervals. Until the $1600th$ acceptance, we set $\gamma_{g,1} = 0.15$, $\gamma_{g,2} = 0.05$, $\gamma_{g,3} = 0.7$, $\gamma_{g,4} = 0.1$, and f_1, f_2 remain unchanged. f_3 is a mixture of normals updated from the sampled set $\theta_0, \theta_1,...$ at regular intervals following the k-harmonic means approach (Hamerly/Elkan, 2002) using the algorithm of Bradly/Fayyad (1998) for initialization. f_4 is identical to f_3 except for the covariance matrices which are 20 times as large. After the $1600th$ acceptance, we replace f_1 by f_3 and set f_2 equal to f_1 but with covariance matrices multiplied by 10, and we keep f_1 and f_2 fixed afterwards. f_3 is obtained again as an estimated mixture of normals, updated regularly, and f_4 is identical to f_3 with covariances 20 times as large. Finally, as parameter estimates we take $\hat{\theta} := \sum_{g=\bar{g}+1}^{G} \theta_g$, where $\bar{g} = 10,000$ is the end of the burn-in period and $G = 20,000$ is the total number of PMCMC iterations.

3.3.3 Simulated Likelihood

Our implementation of the ASIR proceeds as follows for one specific choice of parameters θ:

1. Initialize $k = 1,...,n$ particles by drawing $v_0^k \sim p(v_0|\theta)$ and by setting $\bar{w}_1^k = 1/n$ for $k = 1,...,n$.

2. For $t = 1,...,T-1$:

    (a) $\forall k = 1,...,n$ compute $\pi_{t+1}^k = p(y_{t+1}|\check{v}_{t+1}^k)\bar{w}_t^k$ where $\check{v}_{t+1}^k = \mathbb{E}\left[v_{t+1}|v_t^k\right]$,

    (b) resample (by stratified sampling) the particles so that particle k is chosen with probability $\bar{\pi}_{t+1}^k = \pi_{t+1}^k/\sum_i^n \pi_{t+1}^i$, and note that each v_t^k is redefined by this resampling,

    (c) $\forall k = 1,...,n$ draw $v_{t+1}^k \sim p(v_{t+1}|v_t^k,\theta)$,

    (d) $\forall k = 1,...,n$ compute $w_{t+1}^k = \frac{p(y_{t+1}|v_{t+1}^k)}{p(y_{t+1}|v_t^k)}$ and $\bar{w}_{t+1}^k = w_{t+1}^k/\sum_i^n w_{t+1}^i$.

Note that in step (c), we do not condition on y_{t+1}, which has worked well in the implementation of Johannes et al. (2009) for a univariate stochastic volatility model. It works well in our multivariate framework, too. Pitt et al. (2010) write down the simulated likelihood

$$\hat{p}(y_{1:T}|\theta) = \hat{p}(y_1)\prod_{t=2}^{T}\hat{p}(y_t|y_{1:t-1},\theta) \tag{9}$$

where

$$\hat{p}(y_t|y_{1:t-1},\theta) = \left(\frac{1}{n}\sum_{k=1}^{n}w_t^k\right)\left(\frac{1}{n}\sum_{k=1}^{n}q_t^k\right) \tag{10}$$

and show that this is an unbiased estimator of the likelihood.

3.4 Forecasting

For assessing the quality of forecasts, it is essential to carefully specify the information set used for making each forecast. To do so, we differentiate between the cross-sectional dimension and the time dimension. We further distinguish between knowledge on latent volatility states and on model parameter estimates. We assign indices $i = 1, ..., S$ to all banks in the sample where the indexing reflects the ranking of banks with respect to their size (total assets). Let S be even. We define two forecast samples, $F_1 := \{1, 3, 5, ..., S-1\}$ and $F_2 := \{2, 4, ..., S\}$. We define two learning samples, $L_1 := F_2$ and $L_2 := F_1$. When making forecasts for banks in F_1 we use data from L_1, and when making forecasts for banks in F_2, we use data from L_2.

With respect to the cross-sectional dimension, we treat the banks in the forecast sample as if we had no information about their stock prices. That means, both latent volatility states and model parameters are estimated without knowledge about their returns. With respect to the time dimension, for making a forecast for a specific time interval, we collect latent volatility states estimated with data available up to the last day before the beginning of this time interval. For making reasonable volatility estimates, good estimates of the parameters of the stochastic volatility models are essential. Since pillar 3 information has not been available before 2009, we use estimates of the parameters taken from estimation of the models over the complete sample period January 2009 until November 2011.

We check the quality of volatility forecasts for daily volatility over monthly (20 trading day) intervals by comparing them to empirical volatility for all banks in the forecast sample. We need an estimate of risk factor latent volatility and for idiosyncratic volatility of each bank in the forecast sample as starting values of the latent states in the forecasting simulation. We use the filtered volatility states of the factor model and the average filtered idiosyncratic volatility state across all banks in the *learning* sample, both taken from the last day before the beginning of the estimation period.

We denote by $\hat{\theta}^{(f)}$ the set of parameter estimates for the factor model, by $\hat{\theta}^{(L_1)}$ the set of parameter estimates for the unrestricted stock model obtained using L_1 and by $\hat{\theta}^{(L_2)}$ the set of parameter estimates for the unrestricted stock model using L_2. $\hat{\theta}^{(L_1, \text{restricted})}$ and $\hat{\theta}^{(L_2, \text{restricted})}$ are the estimates obtained under the restriction $\xi_j = 0$, $j = 1, ..., F$.

Then, we filter the factor model and the stock models and define estimates of the latent risk factor logvariances vector for day t as follows:

$$\hat{x}_t := \sum_{k=1}^{n} \bar{w}_t^k \left(x_t^k | \theta^{(f)} \right), \tag{11}$$

where the $\left(x_t^k | \theta^{(f)} \right)$ are observed after completion of step 2d of the ASIR for day t. Similarly, we define $\hat{v}_t^{(L_h)} | \mathcal{F} := \frac{2}{S} \sum_{i \in L_h} \sum_{k=1}^{n} \bar{w}_t^k \left(v_{t,i}^k | \mathcal{F} \right)$ where $h \in \{1, 2\}$ and $\mathcal{F} \in \left\{ \hat{\theta}^{(L_1)}, \hat{\theta}^{(L_2)}, \hat{\theta}^{(L_1, \text{restricted})}, \hat{\theta}^{(L_2, \text{restricted})} \right\}$. Now let t be the first day of a forecast interval of size 20 days. We first simulate (M times) the risk factor volatilities, using the estimated risk factor latent logvariances from $t-1$ as starting values, by sequential updating them according to (2) over 20 trading days. At the same time, we simulate the risk factor returns using the simulated logvariances and (1). We simulate univariate idiosyncratic volatility states according to (4), using the average idiosyncratic volatility state

14

of the banks in the learning sample as starting value, and use these states for all banks in the forecast sample. Then, using the M generated paths of the risk factors and of the idiosyncratic volatility, we simulate bank stock returns, using the risk-factor data and bank-specific exposure data for the forecast sample. Simulation follows (3). Here we use the average mean logvariance parameter estimate $\hat{\mu}_v := \sum_{i=1}^{S} \hat{\mu}_{v,i}$ across all banks in the *learning* sample for sequential simulation of (2). We use $u = 1, ..., M$ simulations and obtain the sets of simulated stock returns $\left\{ y_{\tau,i}^{(u)} | \hat{x}_{t-1}, \hat{v}_{t-1}^{L_2}, \hat{\theta}^{(L_2)} \right\}_{\tau=t,...,t+20; i \in L_1; u=1,...,M}, \left\{ y_{\tau,i}^{(u)} | \hat{x}_{t-1}, \hat{v}_{t-1}^{L_1}, \hat{\theta}^{(L_1)} \right\}_{\tau=t,...,t+20; i \in L_2; u=1,...,M}$ and another pair of such sets with $\hat{\theta}^{(L_1)}, \hat{\theta}^{(L_2)}$ replaced by $\hat{\theta}^{(L_1, \text{restricted})}, \hat{\theta}^{(L_2, \text{restricted})}$. We calculate the volatility forecast associated to the end of day $t-1$ and bank i as

$$\hat{\sigma}_{t-1,i}^2 | \mathcal{F} := \frac{1}{M} \sum_{u=1}^{M} \left(\frac{1}{20-1} \sum_{\tau=t}^{t+19} \left((y_{\tau,i}^{(u)} | \hat{x}_{t-1}, \hat{v}_{t-1}^{L_h}, \mathcal{F}) - \frac{1}{20} \sum_{\tau=t}^{t+19} (y_{\tau,i}^{(u)} | \hat{x}_{t-1}, \hat{v}_{t-1}^{L_h}, \mathcal{F}) \right)^2 \right),$$

where the empirical counterpart is

$$\hat{s}_{t-1,i}^2 := \frac{1}{20-1} \sum_{\tau=t}^{t+19} \left(y_{\tau,i} - \frac{1}{20} \sum_{\tau=t}^{t+19} y_{\tau,i} \right)^2$$

and where $h = 2$ when $i \in F_1$ and $h = 1$ when $i \in F_2$. We assess the quality of forecasts by calculating the unrestricted bias

$$\sum_{i \in F_1} \left(\left(\hat{\sigma}_{t-1,i}^2 | \hat{\theta}^{(L_2)} \right) - \hat{s}_{t-1,i}^2 \right) + \sum_{i \in F_2} \left(\left(\hat{\sigma}_{t-1,i}^2 | \hat{\theta}^{(L_1)} \right) - \hat{s}_{t-1,i}^2 \right)$$

and the unrestricted root mean squared error

$$\sum_{i \in F_1} \left(\left(\hat{\sigma}_{t-1,i}^2 | \hat{\theta}^{(L_2)} \right) - \hat{s}_{t-1,i}^2 \right)^2 + \sum_{i \in F_2} \left(\left(\hat{\sigma}_{t-1,i}^2 | \hat{\theta}^{(L_1)} \right) - \hat{s}_{t-1,i}^2 \right)^2$$

as well as the restricted counterparts where $\hat{\theta}^{(L_1)}, \hat{\theta}^{(L_2)}$ are replaced by $\hat{\theta}^{(L_1, \text{restricted})}, \hat{\theta}^{(L_2, \text{restricted})}$.

4 Data

4.1 Bank Samples

We select the twelve largest (measured by by total assets, at end of 2008) European banks that are listed on a stock exchange and for which pillar 3 reports are available from their website. Most banks start reporting for the 2008 financial year, thus our sample is January 2009 until October 2011. This provides us with a large group of important banks, for which it holds that they are all subject to the same set of disclosure rules. We group these banks by total assets and group together the largest, the third largest, ..., the eleventh-largest bank into a subsample that we denote by "1,3,...". This

subsample has six members. We group together the remaining six banks into another subsample that we denote by "2,4,...". When forecasting for the banks in sample "2,4,.." we use sample "1,3,..." as learning sample, and vice versa. This procedure ensures that no size effect interferes with the distinction between learning and forecasting samples. The banks in our sample are Royal Bank of Scotland (RBS), Deutsche Bank (DB), BNP Paribas (BNP), Barclays (BARC), HSBC, Crédit Agricole (CRAG), UBS, Société Générale (SG), ING, Santander Central Hispano (SAN), Lloyds (LLYD), Credit Suisse (CS). Bank stock returns are daily logreturns, which are demeaned and multiplied by 100, from Thomson Reuters Datastream.

4.2   Factor Returns

We use the following five factors, which we believe are the most important systematic risk drivers for equity returns of European large banks. (1) As corporate credit risk factor, we use the average of sector-specific 5 year - maturity CDS spread indices, as published by Thomson Reuters Datastream, for the following sectors: Auto, Basic Resources, Chemicals, Construction and Materials, Food and Beverages, Industrial, Media, Oil & Gas, Retail, Technology, Telecommunications, Travel and Leisure, Utilities. We do not include spreads for companies in the financial sector. (2) As sovereign credit risk factor, we use a 5 year - maturity CDS index, as published by Thomson Reuters Datastream, for EU sovereigns. (3) For the equity market, we use sector specific MSCI total return stock indices for companies in the European Monetary Union. We translate all of the following index series so that their value is 100 on December 31, 1999: Energy, Materials, Consumer Discretionary, Consumer Staples, Healthcare, IT, Services, Utilities. We calculate a nonfinancial stock index by averaging these normalized indices. (3) As interest rate, we use the yield of the 10 year German Bundesanleihe. (4) For mortgage backed security risk, we use the prices of the iShares exchange traded fund on Barclay's mortgage backed security index. Since this is a price index, we use a total return ETF time series. Since we are interested in factor risk, we calculate daily log-changes to these risk factors, and we demean them and multiply them by 100.

4.3   Exposure Data

We retrieve exposure data from the pillar 3 reports of all banks in our sample. (1) Exposure towards corporate credit risk is measured as total credit exposure in corporate clients with non-investment grade ratings. Since banks use different rating scales, and we use exposure data from the standardized approach and from IRB approaches, we define as non-investment grade any rating where the associated one-year PD is greater than 1%. For calibration, bank-specific disclosures on PDs are used for ratings in the IRB approaches, and Moody's Investors Service (2011a) is used for ratings in the standardized approach. The reason for concentrating on lower quality exposure only is that the sensitivity of the fair value of loans with lower grade ratings towards changes in credit quality is much stronger than for loans with higher grade ratings. (2) Exposure towards sovereign credit risk is measured as total credit exposure in sovereign borrowers with a rating that is associated to a one-year PD of more than 0.05%. For calibration, Moody's Investors Service (2011b) is used. As a measure of exposure to (3) equity risk

respectively (4) interest rate risk, we use the end-of-year Value at Risk of equity respectively interest rate risk in the trading book.[3] We use VaR for a 1 trading day interval at the 99% confidence level. When the trading horizon and/or the confidence level for a bank's VaR disclosures is different, we adjust the disclosed VaR to match our requirements. We do this by assuming a normal distribution for the profit/loss. We add the exposure to equity risk in the banking book to the equity risk exposure. Since the VaR measure is structurally different from the banking book equity exposure data, we need to scale one of them. We choose to multiply VaR by 1000. (5) As a measure of exposure to mortgage backed security risk, we use the sum of securitization exposures in noninvestment grade rating classes. Exposures in the standardized approach are included when it can be observed in the report that their risk weight is above 100%. The exposure data obtained by this procedure is divided by total regulatory capital to adjust for size. Finally, for each risk factor, we normalize the data by dividing each exposure by the average exposure across all banks and years. That is, for any risk factor, exposure above (below) 1 indicates above (below) average exposure.

5 Results

5.1 Descriptive Statistics

First, we present summary statistics for the time series of factor changes and bank-specific stock returns. The empirical distributions of risk factor changes, as shown in table 1, vary greatly. Sovereign spread changes are highly leptokurtic, which corresponds to high variance in the latent variance process (see table 2). Returns of the MBS index exhibit very low volatility, which is surprising at first sight, but has to be seen in the light of the developments on the MBS market prior to the beginning of the sample period. The majority of losses on the MBS market occured before the beginning of 2009.[4]

	Corporate Spread	Sovereign Spread	Stock Market	Interest Rate	MBS Market
Std. Dev.	0.0272	0.0506	0.0142	0.0187	0.0019
Skewness	-0.30	-1.02	-0.14	-0.17	-0.22
Kurtosis	4.85	20.19	1.76	2.06	1.57

Table 1: Factor Returns Time Series

Table 2 shows characteristics of the return distributions for the twelve banks in the complete sample, ordered by size. [5] There is notable variation in standard deviations: the largest bank is the bank with the riskiest returns, the bank with the lowest standard deviation is HSBC. Some return series exhibit substantial excess kurtosis, which in part results from extreme changes of prices that followed news about doubts regarding banks' financial strength. It is also important to note that the return series exhibit substantial variation in volatility over time. A measure of this variability is parameter σ_v in the

---

[3]See sec. **??**, no. 1. Since in some banks, daily VaR is volatile, an average VaR could be better an indicator for interest rate risk exposure, since this could indicate more accurately whether the bank's strategy is to hedge interest rate risk or not.

[4]See also sec. **??**.

[5]These banks represent the twelve largest European banks ranked by total assets. On position 11 is HBOS, which is owned by Lloyds, and therefore position 11 is chosen for Lloyds.
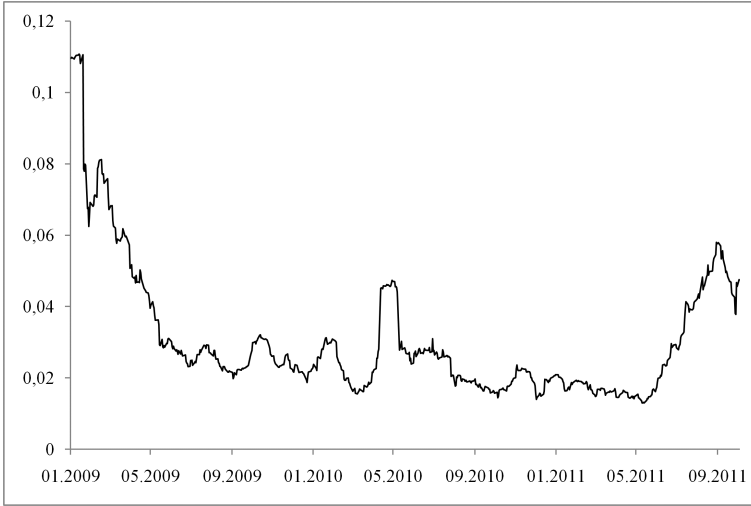
Figure 1: Rolling 20-day window daily bank stock return volatility

stock model (see table 6). Figure 1 shows the average of the rolling window daily standard deviation of all 12 bank stock return series for a window size of 20 days over the sample period January 2009 until October 2011.

	Std. Dev.	Skewness	Kurtosis
RBS	0.0629	-9.89	195.82
DB	0.0322	0.49	5.84
BNP	0.0345	0.55	6.94
BARC	0.0488	1.35	28.46
HSBC	0.0240	-0.44	11.44
CRAG	0.0345	0.21	4.00
UBS	0.0311	-0.18	3.96
SG	0.0380	-0.20	6.83
ING	0.0426	0.02	5.98
SAN	0.0277	0.46	6.66
LLYD	0.0567	-3.58	59.87
CreditSuisse	0.0284	0.21	3.65

Table 2: Stock Returns Time Series

In table 3 we report the exposure measures for all factors and all banks in the sample. Exposure measures are averaged over the complete sample period. The numbers are normalized such that average exposure towards one factor across all banks is one. Since exposures are adjusted for size, they might be used to obtain some insight into the relative sensitivity of banks towards each separate factor. This insight requires, however, that the exposure measures are comparable. This data cannot be used to draw immediate conclusions with respect to total risk for individual banks, because volatility varies across factors, and factors are not independent. The model that we apply is a framework capable of extracting knowledge from the joint evolution of factors and the bank-specific exposures to these factors.

	Corporate Spread	Sovereign Spread	Stock Market	Interest Rate	MBS Market
RBS	0.79	0.13	0.43	0.63	0.52
DB	1.45	0.58	2.30	2.74	1.64
BNP	1.12	1.05	1.52	0.63	0.37
Barclays	0.61	0.58	0.66	0.84	3.67
HSBC	0.71	2.93	0.27	0.36	0.38
CreditAgricole	0.80	0.61	1.21	0.90	0.30
UBS	0.82	2.15	1.51	1.02	0.68
SocieteGenerale	1.23	1.84	1.45	0.51	1.81
ING	1.76	1.14	0.33	0.69	0.76
Santander	0.91	0.14	0.37	0.38	0.21
Lloyds	1.01	0.03	0.09	0.30	0.96
CreditSuisse	0.85	0.70	2.04	3.16	0.40

Table 3: Factor Sensitivities

5.2 Model Parameters

Table 4 contains estimates of the parameters in (2). μ_x measures the mean towards which latent log-variance reverts, and ϕ_x is the speed of reversion.[6] There are remarkable differences among the estimates for different factors. The μ_x estimates generally reflect the average variances of the factor return series. Sovereign spread changes exhibit remarkably high volatility of volatility, which corresponds to high excess kurtosis, and which certainly reflects severe uncertainty of changing degree regarding the soundness of various European sovereign borrowers during the estimation period. We use the mean of the discrete posterior distributions of parameters as estimates, and we the standard deviation of these distributions in parentheses. Hypothesis tests where estimates are tested on being different from zero are not meaningful for the parameters in table 4. From inspection of standard errors we can, however, infer that we can have reasonable confidence in the estimates for most parameters, when comparing the standard deviations to the scale on which parameters are defined (μ_x is defined on the real line, ϕ_x is between 0 and 1, and σ_x^2 is positive and real). Meanwhile, some of the σ_x^2 distributions exhibit substantial dispersion around the mean.

	μ_x	ϕ_x	σ_x^2
Corporate Spread	0.8725	0.9768	0.0079
	(0.1455)	(0.0086)	(0.0098)
Sovereign Spread	0.6402	0.0196	0.9094
	(0.0837)	(0.0107)	(0.1841)
Stock Market	-0.0251	0.5856	0.0942
	(0.0649)	(0.0938)	(0.059)
Interest Rate	0.1656	0.8747	0.0338
	(0.1332)	(0.1065)	(0.03)
MBS market	-1.7023	0.5818	0.2374
	(0.0956)	(0.0672)	(0.1176)

Table 4: Factor Evolution Parameters

---

[6]Note that x is log-variance, hence negative values indicate that the corresponding variance level is below 1.

Estimates of elements in matrix R in (1) and shown in table 5.

	CORP	SOV	EQU	INT	MBS
Corporate Spread	1	0	0	0	0
Sovereign Spread	0.7187***	1	0	0	0
	(0.0882)				
Stock Market	-0.7937***	-0.0163	1	0	0
	(0.0704)	(0.0623)			
Interest Rate	-0.7217***	-0.1116*	0.4916***	1	0
	(0.0692)	(0.0598)	(0.0673)		
MBS market	0.1047**	0.0973	-0.1521**	-0.3061***	1
	(0.0505)	(0.0661)	(0.0803)	(0.0629)	

Table 5: Decomposition of Factor Correlation Matrix

Table 6 shows estimates for the parameters in (4) and for ρ. Variation in the bank-specific means in equation (4) correspond to the variation of return standard deviations: correlation between estimated standard deviations and means, estimated without/with exposure data, is 82.3%/60.8%. Estimated mean reversion speed is similar to estimates from stochastic volatility models for stock indices, see e.g. Chib et al. (2002). The estimate for ρ is negative, which is evidence of a leverage effect in the sense that negative stock returns are likely to be accompanied by volatility increases. The estimates are close to the those from Omori et al. (2007) who find a leverage effect of -0.36 for index returns on the Japanese stock market.

	exposure data			exposure data	
	no	yes		no	yes
μ_v for			μ_v for		
RBS	4.6877	3.4850	DB	2.5561	3.4996
	(1.4612)	(0.7064)		(0.8516)	(0.2488)
BNP	3.1951	2.8649	BARC	2.9849	3.5391
	(0.7752)	(0.5162)		(0.436)	(0.398)
HSBC	2.8391	2.2082	CRAG	2.6435	2.6180
	(0.1922)	(0.6264)		(0.9551)	(0.2975)
UBS	2.7853	2.8618	SG	3.2387	2.6534
	(0.6696)	(0.3735)		(1.3562)	(0.2685)
ING	2.7586	2.8280	SAN	2.1809	3.1234
	(0.4652)	(0.4534)		(0.6675)	(0.2055)
LLYD	3.5665	3.4739	CS	2.6429	3.1242
	(0.689)	(0.6121)		(0.3611)	(0.3074)
ϕ_v	0.9935	0.9896	ϕ_v	0.9896	0.9935
	(0.0039)	(0.0027)		(0.0068)	(0.0024)
σ_v	0.0668	0.0778	σ_v	0.0578	0.0476
	(0.0154)	(0.0135)		(0.0138)	(0.0069)
ρ	-0.3180	-0.2674	ρ	-0.3030	-0.5466
	(0.1849)	(0.2087)		(0.2803)	(0.0916)

Table 6: Stock Evolution Parameters

Next we turn to the sensitivity estimates. Recall that we hypothesized that bank stock returns are driven by risk factors, and that the higher a bank's exposure to a risk factor, the stronger the impact. We focus on two criteria when analyzing the estimates. First, we test whether the sensitivity estimates for factors and for factor-exposure interactions are significantly different from zero. For each parameter, we consider the sampled posterior distribution and count all realizations whose sign is the opposite of the sign of the estimate. We multiply this number by two and divide it by the total number of realizations, which is 10,000. That way, we obtain a p-value which is a simulated equivalent to an asymptotic p-value. When the p-value is smaller than 1%/5%/10%, we indicate this with three/two/one asterisk(s) in table 7. Second, we observe whether the estimates from both samples - banks 1,3,... versus banks 2,4,... - are of the same sign. If this is not the case for both sensitivity parameters for factors and for factor-exposure interactions, we must conclude that our samples are too small to obtain reasonable estimates for the sensitivity parameters. If we obtain estimates of comparable sign and magnitude from both samples for at least either of the two - factor sensitivities or factor-exposure sensitivities - we will be able to draw conclusions with respect to the hypotheses stated in section 2.

	learning from 1,3,...		learning from 2,4,...	
	exposure data		exposure data	
	no	yes	no	yes
β				
Corporate Spread	-0.1479***	-0.1068***	-0.1772***	-0.2744***
	(0.0298)	(0.0392)	(0.0226)	(0.0721)
Sovereign Spread	-0.0107	-0.0097	-0.0182***	0.0094
	(0.0101)	(0.0221)	(0.0073)	(0.0144)
Stock Market	1.0278***	0.8633***	1.0385***	0.9548***
	(0.0534)	(0.0763)	(0.0435)	(0.0699)
Interest Rate	0.0765**	-0.0106	0.1246***	0.259***
	(0.0359)	(0.0554)	(0.0269)	(0.0366)
MBS Market	-0.4321**	-0.6012*	-0.2988*	-0.3458***
	(0.1892)	(0.2711)	(0.148)	(0.1509)
ξ				
Corporate Spread		-0.0718		0.1010
		(0.0586)		(0.0698)
Sovereign Spread		-0.0001		-0.0403***
		(0.0233)		(0.0172)
Stock Market		0.2264***		0.0637
		(0.0639)		(0.0464)
Interest Rate		0.0811**		-0.0871***
		(0.0405)		(0.0177)
MBS Market		0.1908		-0.0206
		(0.2012)		(0.0729)

Table 7: Sensitivity Parameters

For four out of five factors, we obtain estimates which are significantly different from zero and of the same sign (in parentheses) for both samples: corporate credit spread changes (-), nonfinancial stock market returns (+), interest rate changes (+) and returns on the market for mortgage-backed securities

(-). Evidence of a significant (negative) relationship between sovereign credit spread changes and bank stock returns are observed only for the sample of banks 2,4,...; however, for the other sample the sign and the order of magnitude of the estimate is the same. For the first four factors, the estimated sensitivities are hardly surprising. The sensitivity for the MBS market is negative, and some comments on this are in order. When MBS drop in value due to an impairment of the assets backing the issues, so does the equity value of banks that have invested into such assets. When MBS drop in value due to an increase of the interest rate, this increase itself has an effect on an bank's equity value. The direction and magnitude of this effect depend on the maturity structure of the bank's assets and liabilities. Since we observe a positive relationship between interest rate changes and bank stock prices (the corresponding β estimate is positive), we could conclude that a negative β estimate for MBS securities is plausible. However, since the β estimate for MBS securities should only measure the MBS effect on bank returns that is orthogonal to the interest rate effect, a puzzle remains. One possible explanation is that whenever banks are obliged to stand in for guarantees they have given to ensure liquidity of special purpose vehicles (SPV), this will be detrimental to their equity value since these cash injections are unlikely to be repaid. However, these injections will have a positive effect on the value of debt instruments issued by the SPV. This would result in a negative relationship between bank stock returns and returns to MBS securities. Figure 2 shows the posterior distributions of the β estimates taken from sample 2,4,... .
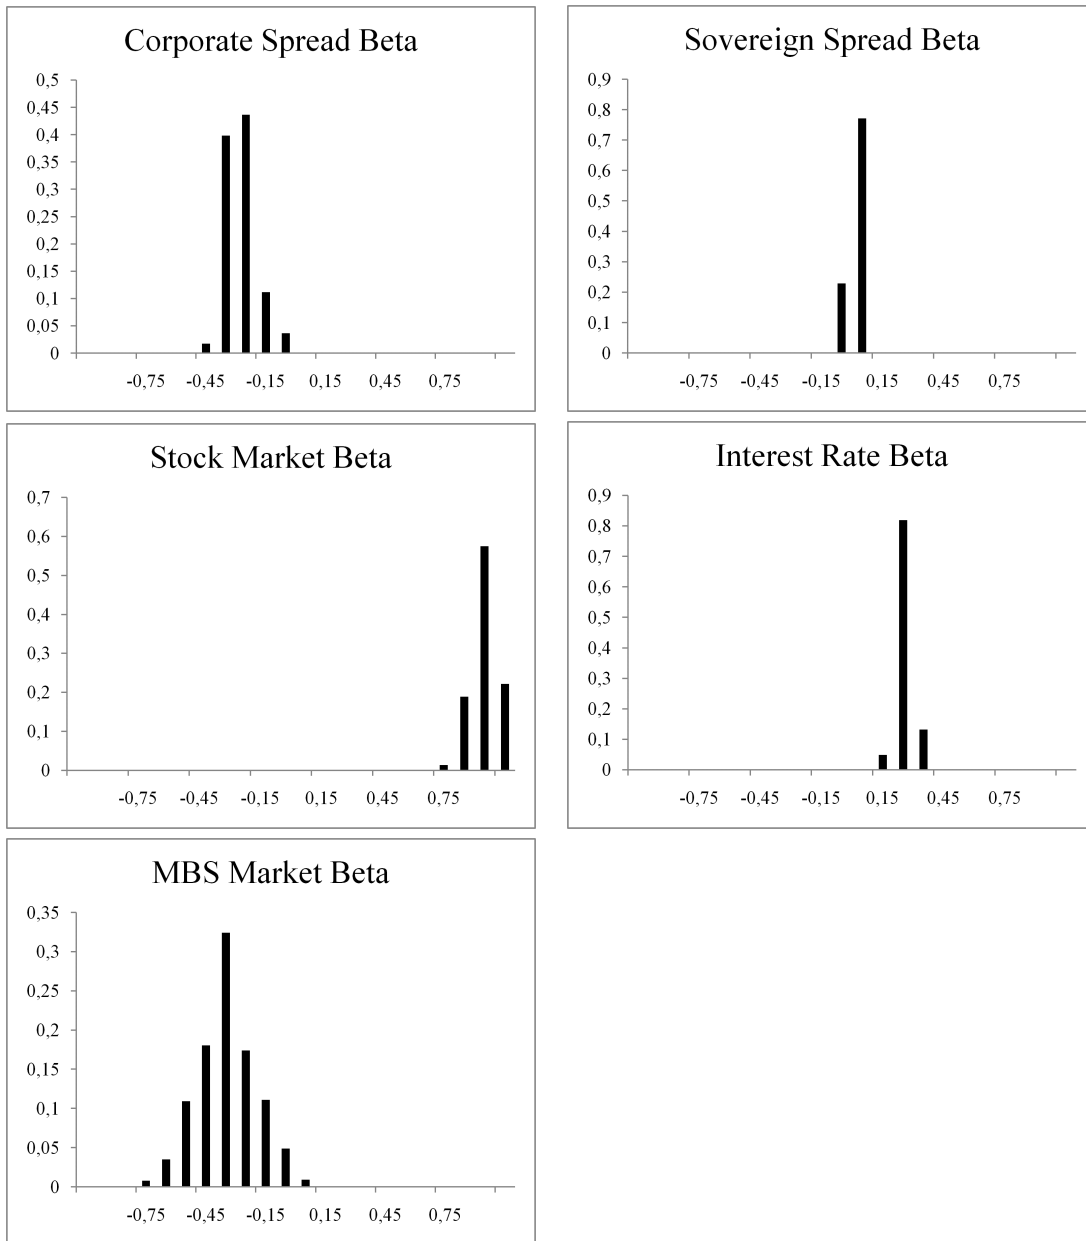
Figure 2: Distribution of Sensitivity Estimates

Things are not as straightforward when considering factor-exposure interactions. There is no factor for which the interaction sensitivity parameter is significantly different from zero and of the same sign in both samples. Thus, we can provide no convincing evidence in support of any of the hypotheses H1 through H5. We can see from table 7 that for sovereign spreads and for nonfinancial stock returns, the coefficient is significantly different from zero for one of the two samples and of the same sign in both samples. This indicates that the exposure data obtained from pillar 3 reports can contribute to some, albeit limited, extent to a careful assessment of an individual bank's sensitivity with respect to these two risk factors, providing limited support for hypothesis H3. The magnitude of this contribution will

be assessed in section 5.3. The sensitivity coefficients for the interest rate exposure interactions are significantly different from zero in both samples, but they are of opposite signs. We could, at first sight, conclude that interest rate VaR of an arbitrary, nonlisted bank is of no value to an analyst who wishes to assess its vulnerability to interest rate changes. At second glance, it is obvious that the direction of the equity value reaction to interest rate changes is not known ex ante, since it depends on the assets' and liabilities' maturity structure. Our results show that banks of the same calibre vary considerably with respect to their net position, otherwise we would not obtain this result. We conclude that more decent (and costly) an analysis of interest rate exposure data is required to draw conclusions with respect to hypothesis H5. In sum, we conclude that risk exposure data is unlikely to substantially improve volatility forecasts. Figure 3 shows the posterior distributions of the ξ estimates for sample 2,4,... .
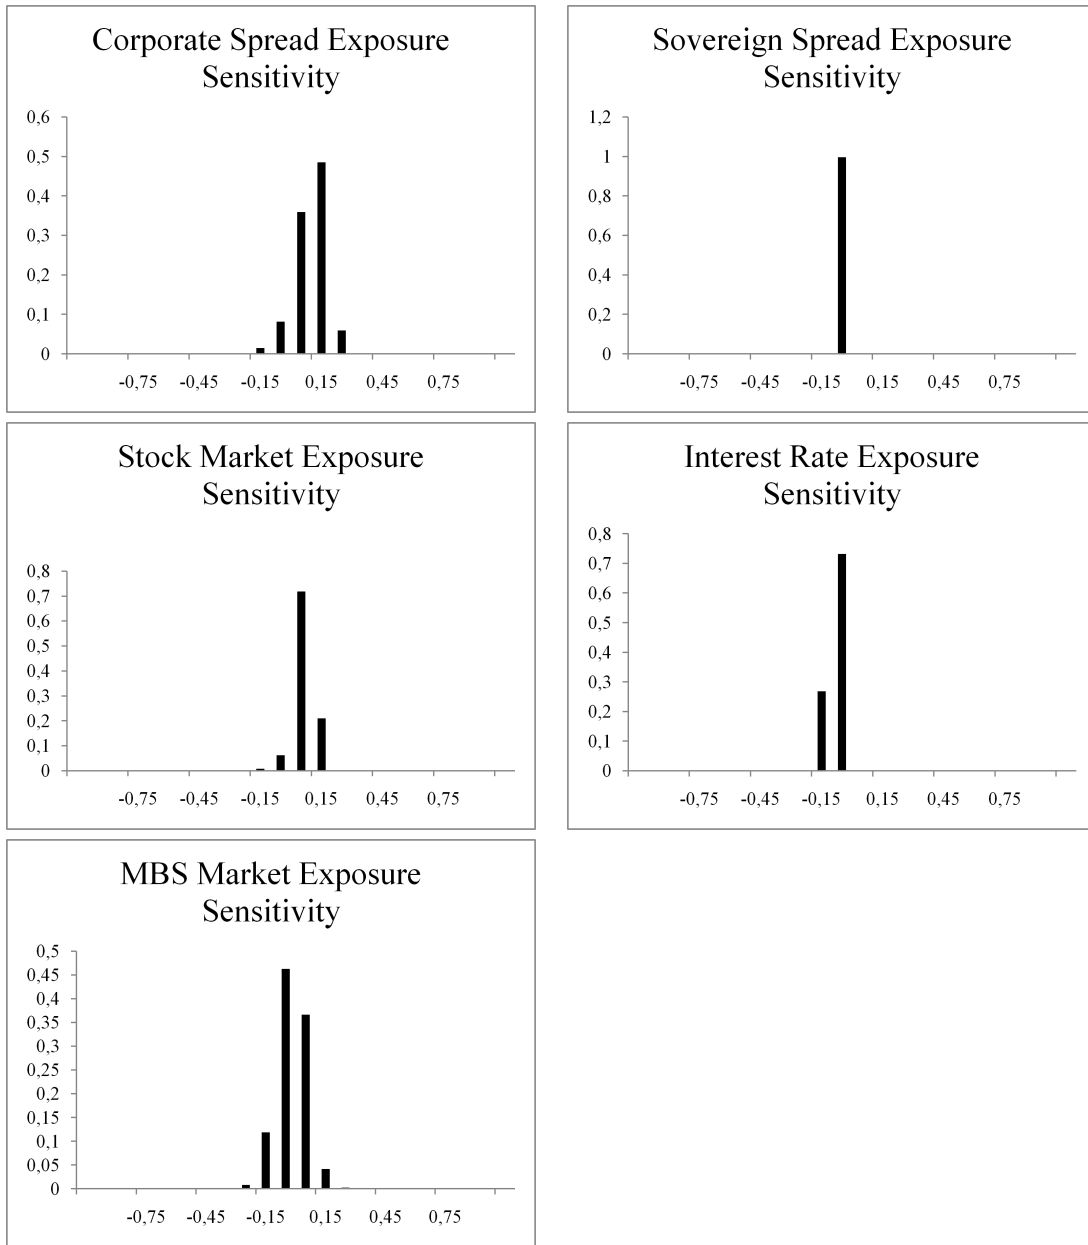
Figure 3: Distribution of Exposure-Sensitivity Interaction Parameters

5.3 Forecasting Results

Table 8 presents average bias and root mean squared error from the out-of-sample forecasts of daily variance over intervals of 20 trading days. We have predictions for 33 of such intervals. The left column contains results for forecasts that have been made without taking into account bank-specific information, i.e. where the forecasts are identical for all six banks. The right column contains results for forecasts which take into account bank-specific information, i.e. the right hand column reflects the

attempt to obtain variance forecasts for individual banks where variation between banks is obtained solely from pillar 3 data. Forecasts are made for two samples, the sample containing the largest, third-largest, ..., and eleventh-largest bank and for the sample containing the second-largest, fourth-largest,..., and twelfth-largest bank. Results are presented separately for both groups. The forecast errors are then aggregated and also reported across the two samples. Generally, bias is positive for both samples. Variance was high especially at the beginning of the sample period. Variance forecasts rely on latent logvariance states at the beginning of the forecast horizon. When variance declines over the sample period on average, then for a given forecast period, the starting value for variance is more likely to be too high than too low. Thus, it is not surprising to see positively biased forecasts.

We concentrate now on H6, which asserts that using pillar 3 information helps to obtain good variance forecasts. In fact, for the first sample, using pillar 3 disclosure data results in better estimates (bias 1.70, RMSE 9.58), both in terms of bias and in terms of RMSE, compared to the uniform forecast that ignores risk exposure data (bias 3.91, RMSE 9.94). For the second sample, the reverse is true (11.20/20.58 versus 8.05/17.19). (Note that return series have been scaled by 100, which is standard in stochastic volatility modeling.) The improvement of forecasts achieved in the first sample is of smaller magnitude than the deteriorioration of forecasts in the second sample. This results in a negative aggregate impact of the use of pillar 3 information on the quality of variance forecasts. Consider the variance forecast for one randomly chosen bank. It is not possible to tell which of the two forecasts we should use. We therefore conclude that there is no robust evidence that the exposure data used in this study is materially useful for predicting equity volatility of banks over 20 trading day periods.[7]

	Exposure data	
	no	yes
predicting for 1,3,...		
Bias	3.91	1.70
RMSE	9.94	9.58
predicting for 2,4,...		
Bias	8.05	11.20
RMSE	17.19	20.58
Combined Samples		
Bias	5.98	6.45
RMSE	14.04	16.05

Table 8: Bias and RMSE for Variance Forecasting

Instead of forecasting variance, it is also possible to forecast a ranking of banks according to their risk using exposure data. For each 20-trading-day interval over which we have calculated a variance forecast, we rank the six banks in forecast according to the forecasts. Then we compare this ranking to the ranking that we obtain ex post upon observing the actual variances. We count the number of ranks that have been predicted correctly in each month. We do this separately for sample 1,3,... and for sample 2,4,... . We count the months for which we have predicted correctly 0/1/.../6 ranks and divide these numbers by the number of periods. We report the resulting frequencies in table 9.

---

[7]We repeated the forecasting procedure for 120 trading day horizons. The magnitude of errors change, but the relative magnitudes do not change, i.e. the same conclusion is drawn from 120 trading day horizons..

To understand the numbers in this table, take the row for 2 recontres as an example. For sample 1,3,..., we report a rounded frequency of 0.3. That means in $0.3 * 33 = 10$ periods, two rank has been correctly predicted. In the second column, we provide the probabilities of obtaining $0/1/.../6$ correct rank predictions if these were made at random (the probabilities are the solution to the *rencontre* - problem, where one "rencontre" means that one rank prediction correctly predicts one rank). If the exposure data allows us to make rank predictions that are more accurate than random predictions, then the observed frequency for one rencontre or more will be higher than the probability associated to this event given random rank predictions. We observe that model-based rank predictions for sample 1,3,... are better than random ranks. This corresponds to the reduction in prediction errors for this sample. However, we observe that also for sample 2,4,... using pillar 3 information provides better rank predictions than random rankings. From this we conclude that the exposure data we chose for this study does provide information about the relative sensitivity of banks with respect to the chosen risk factors. We have argued above that one important purpose of risk disclosures is helping market participants and regulators to estimate the magnitude of bank-specific bank specific equity risk. Our aggregate evidence to date indicates that pillar 3 exposure data is not accurate enough to be of material relevance for this task.

# Rencontres	random ranking	prediction for 1,3,...	prediction for 2,4,...
0	0.79	0.33	0.61
1	0.17	0.27	0.30
2	0.03	0.30	0.06
3	0.01	0.06	0.03
4	<0.01	0.03	0.00
5	<0.01	0.00	0.00
6	<0.01	0.00	0.00

Table 9: Comparison of Random and Exposure-Driven Rankings

6   Summary and Conclusion

A model is designed to assess the informational value of banks' risk disclosures in their Basel II pillar 3 reports. These disclosures are considered as having informational value when they help to forecast individual banks' future equity volatility. While this is why they are published, previous evidence indicates that the quality of such disclosures is weak, i.e. that risk exposure data might not be comparable between banks. We define disjunct learning and forecast samples of European banks and use information from the learning samples to make volatility forecasts for the banks in the forecast samples. We do not use stock return information from the forecast samples. We generate forecasts using the sensitivity of learning sample bank returns with respect to five risk factors (corporate credit risk, sovereign credit risk, nonfinancial stock returns, interest rates, asset-backed securities) and (local in time) risk factor volatility. These forecasts are identical for all banks in the forecast sample since no bank-specific information is used. Then we generate another set of forecasts which take into account the information that pillar 3 risk disclosures provide on the sensitivity of bank returns with respect to the risk factors. After that, we do make use of stock return information for the forecast sample to

compare the quality of forecasts without / with pillar 3 information. The predictions are calculated in a multivariate stochastic volatility model framework on two levels, which is estimated using particle Markov chain Monte Carlo methods. We find that pillar 3 information is capable of producing a ranking between banks in the forecast sample which is better than a random ranking, but that pillar 3 information does not help making good volatility forecasts. One important purpose of pillar 3 risk exposure disclosures is to help assessing the risk of banks in situations where prices of their equity are not available / reliable. We conclude that the quality of pillar 3 information is not sufficient to fulfil this purpose. It must, however, be noted that the size of our sample (12 banks), which is a result of the heavy computational burden of the model, imposes a limitation on our results.

7 References

Andersen, T.G., Bollerslev, T., Christoffersen, P.F., Diebold, F.X. (2006): Volatility and Correlation Forecasting. Handbook of Economic Forecasting 1, 777-878.

Andrieu, C., Doucet, A., Holenstein, R. (2010): Particle Markov chain Monte Carlo methods. Journal of the Royal Statistical Society B 72, 269-342.

Barth, M.E., Beaver, W.H., Landsman, W.R. (2001): The relevance of the value relevance literature for financial accounting standard setting: another view. Journal of Accounting and Economics 31, 77-104.

Berkowitz, J., O'Brien, J. (2002): How accurate are value-at-risk models at commercial banks? Journal of Finance 57, 1093–1111.

Berkowitz, J., Christoffersen, P.F., Pelletier, D. (2006): Evaluating value-at-risk models with desk-level data. Management Science 57, 2213-2227.

Bank for International Settlements (BIS) (2006): International Convergence of Capital Measurement and Capital Standards, Part 3, Basel Committee on Banking Supervision.

Bradley, P.S., Fayyad, U.M. (1998): Refining Initial Points for K-Means Clustering. Proceedings of the 15th International Conference on Machine Learning, 91- 99.

Carey, M. (2001): Some Evidence on the Consistency of Banks' Internal Credit Ratings. In: Credit Ratings: Methodologies, Rationale and Default Risk, Risk Books, London.

Carey, M., Hrycay, M. (2001): Parameterizing credit risk models with rating data. Journal of Banking and Finance 25, 197–270.

Cassidy, C., Gizycki, M. (1997): Measuring traded market risk: Value-at-Risk and Backtesting Techniques. Discussion Paper, Reserve Bank of Australia, November 1997.

Chib, S., Nardari, F., Shephard, N. (2002): Markov chain Monte Carlo methods for stochastic volatility models. Journal of Econometrics 108, 281-316.

European Parliament and European Council (2006): Directive 2006/48/EC of 14 June 2006 relating to the taking up and pursuit of the business of credit institutions. Official Journal of the European Union 2006, L 177/1 - 197.

Frésard, L., Pérignon, C., Wilhelmsson, A. (2011): The pernicious effects of contaminated data in risk management. Journal of Banking & Finance 35, 2569-2583.

González, F. (2005): Bank regulation and risk-taking incentives: An international comparison of bank risk. Journal of Banking & Finance 29, 1153-1184.

Holthausen, R.W., Watts, R.L. (2001): The relevance of the value-relevance literature for financial accounting standard setting. Journal of Accounting and Economics 31, 3-75.

Jacobson, T., Lindé, J., Roszbach, K. (2006): Internal ratings systems, implied credit risk and the consistency of banks' risk classification policies. Journal of Banking & Finance 30, 1899-1926.

Johannes, M.S., Polson, N., Stroud, J.R. (2009): Optimal Filtering of Jump Diffusions: Extracting Latent States from Asset Prices. The Review of Financial Studies 22, 2559-2599.

Moody's Investors Service (2011a): European Corporate Default and Recovery Rates, 1985 - 1H 2011 (Excel Supplement).

Moody's Investors Service (2011b): Sovereign Default and Recovery Rates, 1983-2010. Special Comment.

Pitt, M.K., Silva, R.S., Giordani, P., Kohn, R. (2010): Auxiliary particle filtering within adaptive Metropolis-Hastings sampling. Working Paper, University of Warwick, June 2010.

Lucas, A. (2001): An evaluation of the Basle guidelines for backtesting banks' internal risk management models. Journal of Money, Credit and Banking 33, 826–846.

Omori, Y., Chib, S., Shephard, N., Nakajima, J. (2007): Stochastic volatility with leverage: Fast and efficient likelihood inference. Journal of Econometrics 140, 425-449.

Pérignon, C., Deng, Z.Y., Wang, Z.J. (2008): Do banks overstate their Value-at-Risk? Journal of Banking & Finance 32, 783-794.

Pérignon, C., Smith, D.R. (2010a): The level and quality of Value-at-Risk disclosure by commercial banks. Journal of Banking & Finance 34, 362-377.

Pérignon/Smith (2010b): Diversification and Value-at-Risk. Journal of Banking & Finance 34, 55-66.

Pitt, M., Shephard, N. (1999): Filtering via Simulation: Auxiliary Particle Filters. Journal of the American Statistical Association 94, 590-599.

Roberts, G.O., Rosenthal, J.S. (2009): Examples of adaptive MCMC. Journal of Computational and Graphical Statistics 18, 349–367.

Wong, W.K. (2008): Backtesting trading risk of commercial banks using expected shortfall. Journal of Banking & Finance 32, 1404–1415.

SFB 649 Discussion Paper Series 2012

For a complete list of Discussion Papers published by the SFB 649,
please visit http://sfb649.wiwi.hu-berlin.de.

001   "HMM in dynamic HAC models" by Wolfgang Karl Härdle, Ostap Okhrin
and Weining Wang, January 2012.

002   "Dynamic Activity Analysis Model Based Win-Win Development
Forecasting Under the Environmental Regulation in China" by Shiyi Chen
and Wolfgang Karl Härdle, January 2012.

003   "A Donsker Theorem for Lévy Measures" by Richard Nickl and Markus
Reiß, January 2012.

004   "Computational Statistics (Journal)" by Wolfgang Karl Härdle, Yuichi Mori
and Jürgen Symanzik, January 2012.

005   "Implementing quotas in university admissions: An experimental
analysis" by Sebastian Braun, Nadja Dwenger, Dorothea Kübler and
Alexander Westkamp, January 2012.

006   "Quantile Regression in Risk Calibration" by Shih-Kang Chao, Wolfgang
Karl Härdle and Weining Wang, January 2012.

007   "Total Work and Gender: Facts and Possible Explanations" by Michael
Burda, Daniel S. Hamermesh and Philippe Weil, February 2012.

008   "Does Basel II Pillar 3 Risk Exposure Data help to Identify Risky Banks?"
by Ralf Sabiwalsky, February 2012.