

Observation Driven Mixed-Measurement Dynamic Factor Models with an Application to Credit Risk

Drew Creal^a, Bernd Schwaab^{b*}, Siem Jan Koopman^{c,e}, André Lucas^{d,e}

^(a) *Booth School of Business, University of Chicago*

^(b) *European Central Bank*

^(c) *Department of Econometrics, VU University Amsterdam*

^(d) *Department of Finance, VU University Amsterdam & Duisenberg School of Finance*

^(e) *Tinbergen Institute, Amsterdam*

February 11, 2011

Abstract

We propose a dynamic factor model for mixed-measurement and mixed-frequency panel data. In this framework time series observations may come from a range of families of parametric distributions, may be observed at different time frequencies, may have missing observations, and may exhibit common dynamics and cross-sectional dependence due to shared exposure to dynamic latent factors. The distinguishing feature of our model is that the likelihood function is known in closed form and need not be obtained by means of simulation, thus enabling straightforward parameter estimation by standard maximum likelihood. We use the new mixed-measurement framework for the signal extraction and forecasting of macro, credit, and loss given default risk conditions for U.S. Moody's-rated firms from January 1982 until March 2010.

Keywords: panel data; loss given default; default risk; dynamic beta density; dynamic ordered probit; dynamic factor model.

JEL classification codes: C32, G32.

*We thank Emanuel Moench (our discussant at the 2011 AFA meetings in Denver), and Michel van der Wel for earlier comments. We are grateful to seminar participants at CREATES at Aarhus, Tinbergen Institute, the University of Chicago Booth, and VU University Amsterdam. Schwaab thanks the C. Willems Stichting for financial support while visiting the University of Chicago, Booth School of Business. Corresponding author: Drew Creal, University of Chicago, Booth School of Business, 5807 S. Woodlawn Ave., Chicago, IL 60637; Tel 773.834.5249, Email: dcreal@chicagobooth.edu. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the European Central Bank or the European System of Central Banks.

1 Introduction

Consider an unbalanced panel time series y_{it} for $i = 1, \dots, N$ and $t = 1, \dots, T$, where each variable can come from a different distribution. Such settings occur naturally in many areas of economics and the need for a joint modeling framework for variables from different distributions with common features has become more apparent recently. For example, to construct an accurate and reliable business cycle indicator, many different measurements of economic activity need to be considered simultaneously. Some of these may be Gaussian distributed (typical macroeconomic time series), whereas others are fat-tailed (such as stock returns), integer or binary (such as NBER recession dates), or categorical (such as consumer confidence that is indicated as low, moderate, high) variables. All of these variables reflect a common exposure to the business cycle, but at the same time each variable needs its own appropriate distributional specification. In this paper, we propose an observation driven dynamic modeling framework for the simultaneous analysis of mixed-measurement time series that are subject to common features.

An additional challenge when multiple time series are available is that the observation frequencies can be different for each time series. Some series are observed every year while other series are observed every quarter or month. In empirical financial studies, daily and intra-daily time series of returns are also readily available for analysis. A simultaneous analysis of time series with different observation frequencies is a challenging task. Different methodologies have been developed for this purpose. For example, Mariano and Murasawa (2003) adopt a state space approach for the construction of a coincident business cycle index using quarterly and monthly data while Ghysels, Santa Clara, and Valkanov (2006) show that the precision in predicting the volatility of financial time series can benefit from a mixed-data sampling analysis applied to intra-daily returns of different frequencies. Our mixed-measurement modeling framework incorporates a mixed-data sampling approach by explicitly formulating a high-frequency time series process and allowing for missing observations in the analysis.

Our main motivation to develop a mixed-measurement, mixed-frequency modeling framework is for the estimation, analysis and forecasting of credit risk. Credit risk analysis has been highly relevant in the aftermath of the 2008 financial crisis. Financial institutions and regulators are specifically trying to assess what is the common variation in firm defaults in

order to correctly assess risk. In our empirical analysis we focus on the systematic variation in cross-sections of macroeconomic data, credit rating transitions, and bond loss rates upon default (also known as loss given default). Our data set exhibits the complications as discussed above. While the number of credit ratings transitions between rating categories is a discrete and even ordered random variable, the macroeconomic variables are modeled as continuous variables, whereas the percentage amounts lost on the principal in case of default are continuous and bounded between zero and one. Some of the macro series are observed quarterly while others are observed monthly. Furthermore, the loss given defaults are only observed if there are defaults, such that we have many missing observations by construction. Finally, all series exhibit some common dynamic features related to the business cycle. Loss rates and defaults both tend to be high during an economic downturn, indicating important systematic covariation across different types of data. The commonalities are captured by latent dynamic factors in our modeling framework. The total data set forms an unbalanced panel with 19,540 rating transition events for 7,505 companies, 1,342 cases of (irregularly spaced) defaults with associated losses given default, and six macroeconomic series of mixed quarterly and monthly frequency. The complicated nature of the data set underlines the flexibility of the observation driven modeling framework as developed in this paper.

When the parameters in the model are estimated, we can use the model to forecast credit risk conditions in the economy and to construct predictive loss densities for portfolios of corporate bonds at different forecasting horizons. The model can therefore be used to stress test current credit portfolios and determine adequate capital buffers using the high percentiles of the simulated portfolio loss distributions. Our modeling framework provides a relatively simple observation driven alternative to the (parameter driven) frailty models of McNeil and Wendin (2007), Koopman, Lucas, and Monteiro (2008), and Duffie, Eckner, Horel, and Saita (2009). In addition, our proposed model allows the identification of three components of credit risk simultaneously: macro, default and rating migration, and loss given default. In earlier work, models have concentrated on defaults only, defaults and ratings, or defaults and macro risk.

The distinguishing feature of our modeling framework is that it is entirely observation driven. Observation driven time series models allow parameters to vary over time as functions of lagged dependent variables and exogenous variables. The parameters are stochastic but perfectly predictable given the past information. In the alternative class of parameter driven models this

perfect predictability is lost; see Cox (1981) for a distinction between the two classes of models. The main advantage of an observation driven approach is that the likelihood is known in closed form. It leads to simple procedures for likelihood evaluation and in particular it avoids the need for simulation based methods to evaluate the likelihood. Observation driven time series models have become popular in the applied statistics and econometrics literature. Typical examples of these models include the generalized autoregressive conditional heteroskedasticity (GARCH) model of Engle (1982) and Bollerslev (1986), the autoregressive conditional duration (ACD) model of Engle and Russell (1998), and the dynamic conditional correlation (DCC) model of Engle (2002). In the same spirit, we develop a panel data model for mixed-frequency observations from different families of parametric distributions which are linked by a small set of latent dynamic factors. The likelihood of this new model exists in closed form and can be maximized in a straightforward way.

A number of well-known methods for the modeling of large panels of time series based on latent dynamic factors have been explored: (i) principal components analysis in an approximate dynamic factor model framework, see e.g. Connor and Korajczyk (1988, 1993), Stock and Watson (2002), Bai (2003), Bai and Ng (2002, 2007); (ii) frequency-domain estimation, see e.g. Sargent and Sims (1977), Geweke (1977), Forni, Hallin, Lippi, and Reichlin (2000, 2005); and (iii) signal extraction using a state space time series analysis, see e.g. Doz, Giannone, and Reichlin (2006), and Jungbacker and Koopman (2008). Compared to the methods of (i) and (ii), our current framework provides an integrated parametric framework for obtaining in-sample estimates and out-of-sample forecasts for the latent factors and other variables in the model. Compared to the methods under (iii), our likelihood is known in closed form, even when the model is (partially) nonlinear and includes non-Gaussian densities. Our modeling framework provides basic and simple procedures for likelihood evaluation and parameter estimation without compromising the flexibility of model formulations that aim to construct effective forecasting distributions.

In Section 2, we introduce observation driven mixed-measurement dynamic factor models. We then provide an application of the new framework in Section 3 to jointly model macroeconomic dynamics, credit rating dynamics, defaults, and losses given default. In Section 4, we use the new model to estimate and forecast time-varying credit risk and loss given default risk factors jointly with macroeconomic variables at a business cycle frequency. Section 5 concludes.

2 Mixed-measurement dynamic factor models

This section introduces the observation driven mixed-measurement dynamic factor model for the modeling of a large unbalanced panel of time series. The methodology for the extraction of the factor and the maximum likelihood estimation of the parameters in the model needs to allow for missing values. Since different time series can be observed at different frequencies and each time series can be observed within different time intervals, missing observations are common place in our analysis.

2.1 Model specification

Consider the $N \times 1$ vector of variables y_t of which N_t elements are observed and $N - N_t$ elements are treated as missing, with $1 \leq N_t \leq N$, at time period t . The measurement density for the i th element of y_t is given by

$$y_{it} \sim p_i(y_{it}|f_t, \mathcal{F}_{t-1}; \psi), \quad \text{for } i = 1, \dots, N, \quad t = 1, \dots, T, \quad (1)$$

where f_t is vector of unobserved factors or time-varying parameters, $\mathcal{F}_t = \{y_1, \dots, y_t\}$ is the set of past and concurrent observations at time t , and ψ is a vector of static unknown parameters. In our mixed-measurement framework, the densities $p_i(y_{it}|f_t, \mathcal{F}_{t-1}; \psi)$ for $i = 1, \dots, N$, can originate from different families of distributions. All distributions however depend upon the same $M \times 1$ vector of common unobserved factors f_t . We assume a factor model structure in which the y_{it} 's at time t are cross-sectionally independent conditional on f_t and on information set \mathcal{F}_{t-1} . We then have

$$\log p(y_t|f_t, \mathcal{F}_{t-1}; \psi) = \sum_{i=1}^N \delta_{it} \log p_i(y_{it}|f_t, \mathcal{F}_{t-1}; \psi), \quad (2)$$

where δ_{it} takes the value one when y_{it} is observed and zero when it is missing. The density in (2) may also depend on a vector of exogenous covariates. We omit this extension here to simplify the notation.

The dynamic factor f_t is modeled as an autoregressive moving average process given by

$$f_{t+1} = \omega + \sum_{i=1}^p A_i s_{t-i+1} + \sum_{j=1}^q B_j f_{t-j+1}, \quad t = 1, \dots, T, \quad (3)$$

where s_1, \dots, s_T is a martingale difference sequence with mean zero, ω is an $M \times 1$ vector of constants and the coefficients A_i and B_j are $M \times M$ parameter matrices for $i = 1, \dots, p$ and $j = 1, \dots, q$. The coefficients can be specified and restricted so that the process f_t is covariance stationary. The unknown static parameters in (1) together with the unknown elements in ω , A_1, \dots, A_p and B_1, \dots, B_q are collected in the static parameter vector ψ . The initial value f_1 is taken as fixed at the unconditional mean of the stationary process f_t .

We follow Creal, Koopman, and Lucas (2010) by setting the innovation s_t in (3) equal to the score of the log-density $p(y_t|f_t, \mathcal{F}_{t-1}; \psi)$ for $t = 1, \dots, T$. In particular, s_t is defined as

$$s_t = S_t \nabla_t, \quad \text{where} \quad \nabla_t = \frac{\partial \log p(y_t|f_t, \mathcal{F}_{t-1}; \psi)}{\partial f_t}, \quad (4)$$

and where S_t is an appropriately chosen scaling matrix. The scaled score s_t in (3) is a function of past observations, factors, and unknown parameters. It follows immediately from the properties of the score that the sequence s_1, \dots, s_t is a martingale difference. The dynamic factors f_t are therefore driven by a sequence of natural innovations.

For particular choices of the measurement density $p(y_t|f_t, \mathcal{F}_{t-1}; \psi)$ and the scaling matrix S_t , Creal, Koopman, and Lucas (2010) show that the modeling framework (2)-(3) reduces to popular models such as the GARCH model of Engle (1982) and Bollerslev (1986), the ACD model of Engle and Russell (1998), the multiplicative error model of Engle and Gallo (2006), as well as other models. For our mixed-measurement model in which we allow for missing observations and for different observation frequencies, we construct the scaling matrix from the eigenvalue-eigenvector decomposition of the Fisher information matrix as given by

$$\mathcal{I}_t = E_{t-1}[\nabla_t \nabla_t'] = E[\nabla_t \nabla_t' | f_t, \mathcal{F}_{t-1}].$$

The eigendecomposition of matrix \mathcal{I}_t is represented by

$$\mathcal{I}_t = U_t \Sigma_t U_t',$$

with the columns of $M \times r$ matrix U_t equal to the eigenvectors of \mathcal{I}_t corresponding to its nonzero eigenvalues, and $r \times r$ diagonal matrix Σ_t containing the nonzero eigenvalues of \mathcal{I}_t . We have implicitly defined r as the rank of \mathcal{I}_t . The scaling matrix is then

$$S_t = U_t \Sigma_t^{-1/2} U_t' \quad (5)$$

which can be regarded as the generalized square root inverse matrix of \mathcal{I}_t . By having S_t based on the Fisher information matrix, the gradient ∇_t is corrected for the local curvature of the measurement density $p(y_t|f_t, \mathcal{F}_{t-1}; \psi)$ at time t . Furthermore, the martingale difference series s_t has a finite, idempotent covariance matrix. For example, when the information matrix is nonsingular, the covariance matrix of s_t equals the identity matrix.

In the mixed-measurement setting with measurement densities specified by (1) and (2), the score vector at time t takes a simple additive form

$$\nabla_t = \sum_{i=1}^N \delta_{it} \nabla_{i,t} = \sum_{i=1}^N \delta_{it} \frac{\partial \log p_i(y_{it}|f_t, \mathcal{F}_{t-1}; \psi)}{\partial f_t}, \quad (6)$$

where δ_{it} takes the value one when y_{it} is observed and zero when it is missing. Similarly, the conditional information matrix is additive

$$\mathcal{I}_t = E_{t-1}[\nabla_t \nabla_t'] = \sum_{i=1}^N \delta_{it} E_{i,t-1}[\nabla_{i,t} \nabla_{i,t}']. \quad (7)$$

Given these results, it is straightforward to compute the scaling matrix in (5).

2.2 Measurement of the factors

The estimation of the factors at time t given past observations $\mathcal{F}_{t-1} = \{y_1, \dots, y_{t-1}\}$, for a given value of ψ , is carried out as a filtering process. At time t , we assume that \mathcal{F}_{t-1} and the paths f_1, \dots, f_t and s_1, \dots, s_{t-1} are given. When observation y_t becomes available, we compute s_t as defined in (4) with scaling matrix (5). Subsequently, we compute f_{t+1} using the recursive equation (3). At time $t+j$, once observation y_{t+j} is available, we can compute s_{t+j} and f_{t+1+j} in the same way for $j = 1, 2, \dots$. In practice, the filtering process is started at $t = 1$ with f_1 being set to some fixed value. The initial value f_1 can also be treated as a part of ψ .

Missing values in data sets are intrinsically handled simply through the specifications of ∇_t and \mathcal{I}_t in (6) and (7), respectively. The variables ∇_t and \mathcal{I}_t enable the computation of s_t . When all entries in y_t are missing, it follows that $s_t = 0$. The computation of f_{t+1} using (3) is not affected further when we have missing values. Hence our modeling framework adapts naturally to missing values.

When the time series panel is unbalanced, missing values appear naturally in the data at the beginning and/or at the end of the time series. Also they appear when time series are observed at different frequencies. The overall time index refers to a time period associated with the highest available frequency in the panel. Time series observed at lower frequencies contain missing values at time points for which no new observations are available. For example, a panel with monthly and quarterly time series adopts a monthly time index and a quarterly time series is arranged by having two monthly missing values after each (quarterly) observation. The precise arrangement depends on whether the variable represents a stock (point in time) or a flow (a quantity over time, or average).

2.3 Maximum likelihood estimation

Observation driven time series models are attractive because the log-likelihood is known in closed form. For a given set of observations y_1, \dots, y_T , the vector of unknown parameters ψ can be estimated by maximizing the log-likelihood function with respect to ψ , that is

$$\hat{\psi} = \arg \max_{\psi} \sum_{t=1}^T \log p(y_t | f_t, \mathcal{F}_{t-1}; \psi), \quad (8)$$

where $p(y_t | f_t, \mathcal{F}_{t-1}; \psi)$ is defined in (2). The evaluation of $\log p(y_t | f_t, \mathcal{F}_{t-1}; \psi)$ is easily incorporated in the filtering process for f_t as described in Section 2.2.

The maximization in (8) can be carried out using a conveniently chosen quasi-Newton optimization method that is based on score information. The score here is defined as the first derivative of the log-likelihood function in (8) with respect to the constant parameter vector ψ . Analytical expressions for the score function can be developed but it typically leads to a collection of complicated equations. In practice, the maximization of the likelihood function is therefore based on numerical derivatives.

Identification of the individual parameters in ψ needs to be considered carefully in factor models. A rotation of the factors by some nonsingular matrix may yield an observationally equivalent model. To make sure that all coefficients in ψ are identified, we impose the restriction $\omega = 0$ in (3) and we restrict the set of factor loadings. In particular, we restrict M rows in the factor loading matrix to have a lower triangular format with ones on the diagonal, and we assume the matrices A_i and B_j of (3) for $i = 1, \dots, p$ and $j = 1, \dots, q$ to be diagonal.

2.4 Forecasting

The forecasting of future observations and factors is straightforward. The forecast f_{T+h} , with $h = 1, 2, \dots, H$, can be obtained by iterating the factor recursion (3) in which the sequence s_{T+1}, \dots, s_{T+H} is treated as a martingale difference. To obtain forecasting expectations of nonlinear functions of the factors, the conditional mean of the predictive distribution needs to be computed by simulation due to Jensen's inequality. Simulating the factors is straightforward given the recursion (3). Simulation is also the appropriate tool if other characteristics of the forecasting distribution are of interest such as percentiles and quantiles.

Forecasting in our modeling framework has several advantages when compared to the two-step forecasting approach in the approximate dynamic factor modeling framework of Stock and Watson (2002). First, forecasting the future observations and factors in our framework does not require the formulation of an auxiliary model. Parameter estimation, signal extraction, and forecasting occurs in a single unified step. In the two-step approach, first, the factors are extracted from a large panel of predictor variables, and, second, the forecasts for the variables of interest are computed via regression with the lagged estimated factors as covariates. Our simultaneous modeling approach is conceptually more straightforward, it retains valid inference that may be lost in a two-step approach, and it ensures that the extracted factors are related to the variables of interest throughout the estimation and forecasting process.

3 An application to macroeconomic and credit risk

The focus on credit risk has increased considerably since the 2007-2008 financial crises in both the professional and academic finance literature. Credit risk is often discussed in terms of the

probability of default (PD) and the loss given default (LGD): PD is the probability that a firm or company goes into default and LGD is the fraction of the capital that is lost in case the firm is in default. It is argued that both PD and LGD are driven by the same underlying risk factors; see the discussions in Altman, Brady, Resti, and Sironi (2003), Allen and Saunders (2004), and Schuerman (2006). The implication is that LGD is expected to be high when PD is expected to be high as well. As a result, the total credit risk profile of a portfolio increases.

In our empirical study, we apply the general modeling framework of Section 2 to investigate the linkages between macroeconomic and credit risk. We analyze firm-level data on defaults and on changes in credit quality to obtain insight into the dynamic relations between LGD and macroeconomic fluctuations. The model for credit quality is based on a dynamic ordered logit distribution and the model for LGD is based on a dynamic beta distribution; see Gupton and Stein (2005) and CreditMetrics (2007) for static versions of our model. The macroeconomic variables are specified as a linear Gaussian dynamic factor model.

3.1 Data

Our available time series panel consists of three groups of variables: macroeconomic, default and loss given default (LGD). The macroeconomic group has six time series (five monthly and one quarterly) that we have obtained from the FRED database at the Federal Reserve Bank of St. Louis. General macroeconomic conditions are reflected by three variables : annual change in industrial production growth (monthly), annual change in the unemployment rate (monthly), and annual change in real GDP (quarterly). These three variables are strongly related to the state of the business cycle and measure the extent of economic activity. General financial market and credit variables are included to account for probability to default conditions as perceived by the market. We include the three monthly variables of credit spread, annual change in stock market log-prices (returns), and stock market volatility. The credit spread is measured as the spread between the yield on Baa rated bonds and treasury bonds, where the ratings are assigned by Moody's. Credit spread movements capture two components of credit risk: changes in the market's perception of the probability of default and the losses given default; changes in the price that the market charges for this type of risk. Particularly the first of these two elements can be relevant for determining default rate dynamics. The stock market variables are

the yearly returns on the S&P 500 index and the volatility of the same index. The volatility is measured by the daily realized volatility computed over the past month. Both variables can be linked to default risk through the structural model of Merton (1974) in which firms with higher asset values or lower asset volatilities are less likely to default. In the aggregate, the dynamics of the two can be approximated by equity returns and equity volatilities given that the average debt-equity proportions of the S&P index constituents are relatively stable over time. The sample period January, 1982 to March, 2010 contains 350 observations. We standardize all six macroeconomic variables for our analysis by subtracting the (time series) mean and dividing by the standard deviation.

The default group of variables contains credit ratings assigned by Moody's reflecting the credit quality of the firm. The rating of firm i at time t is denoted by R_{it} . We map the ratings scale of Moody's on a coarser grid of four ratings. The categories are labeled Investment Grade (IG) containing Moody's rating grades Aaa down to Baa3; double B (BB) containing Ba1–Ba3; single B (B) containing B1–B3; and triple C (CCC) containing Caa1–C3. All companies in the sample can also default, which is marked as a transition to the absorbing category D. We have $R_{it} \in \{\text{IG, BB, B, C, D}\}$. To account for all possibilities including defaults, we have five transitions that a firm can make from its current rating (including staying in its current rating). We therefore keep track of sixteen possible types of rating transitions. Each firm also belongs to one of eleven industry categories which we consolidate into seven: financial, transportation, hotel & media, utilities & energy, industrials, technology, and retail & consumer products. In April 1982 and October 1999 Moody's redefined some of their rating categories. These events cause a large number of rating transitions for some categories in these months. We handle these events in our model by including dummy variables for these two periods.

The loss given default (LGD) measures the fraction of the total exposure that is lost conditional on a firm defaulting. Our sample contains 1342 defaults from which we have 1125 measurements of LGD. The number of observed LGDs can be further separated by industry into 100 financial, 48 transportation, 188 hotel & media, 94 utilities & energy, 359 industrials, 139 technology, and 177 retail & consumer products. The LGD is measured from financial market data using what is known as the market implied LGD. Market implied LGDs are constructed by recording the price of a traded bond just before the default announcement and the market price of the same bond 30 days after the default announcement. The percentage drop in

price then defines the loss fraction or LGD; see McNeil, Frey, and Embrechts (2005) for further details on the different ways to measure LGDs. Missing LGDs in the database are due to the underlying bonds not being traded in the market or to the unavailability of price information on the bonds in the underlying data sources. The LGD variable is a vector with its element representing a default at time t . The dimension K_t of this vector therefore varies over time.

3.2 Model for credit risk

From the above data, our observation vector can be separated into three subvectors $y_t = (y_t^m, y_t^c, y_t^r)'$ containing the macroeconomic, credit rating, and LGD data, respectively. Time-variation in the macroeconomic data at the business cycle frequency is assumed to carry over into the credit ratings and the loss given default rates. Consequently, these data share a common exposure to the dynamic latent risk factors f_t . Our parsimonious model for the measurement densities is given by

$$y_t^m \sim N(\mu_t, \Sigma_m) \quad (9)$$

$$y_{i,t}^c \sim \text{Ordered Logit}(\pi_{ijt}, j \in \{\text{IG, BB, B, C, D}\}), \quad (10)$$

$$y_{k,t}^r \sim \text{Beta}(a_{kt}, b_{kt}), \quad k = 1, \dots, K_t, \quad (11)$$

where the 6×1 mean vector $\mu_t = \mu(f_t)$ is a function of factor f_t , the 6×6 variance matrix Σ_m is fixed over time, the probability π_{ijt} concerns firm i and its transition from rating R_{it} to rating j during period t for $j = 1, \dots, 5$ with $\pi_{ijt} = \pi_{ij}(f_t)$ as a function of f_t . The k th element of y_t^r is the LGD fraction for the k th default at time t and is modeled by a beta distribution with coefficients $a_{kt} = a_k(f_t)$ and $b_{kt} = b_k(f_t)$ both functions of f_t . Further details and discussions related to our measurement densities are presented below.

Macro model

The macroeconomic vector y_t^m is specified by the multivariate normal distribution in (9). Let $y_t^{sm} = \tilde{S}_t y_t^m$ be the subvector of y_t^m that is observed at time t , where \tilde{S}_t is an appropriately sized selection matrix. The log-likelihood contribution at time t for y_t^m is

$$\text{const} - 0.5 \log |\tilde{S}_t \Sigma_m \tilde{S}_t'| - 0.5 \left(\tilde{S}_t (y_t^m - \mu_t) \right)' \left(\tilde{S}_t \Sigma_m \tilde{S}_t' \right)^{-1} \left(\tilde{S}_t (y_t^m - \mu_t) \right), \quad (12)$$

with the 6×1 time-varying mean vector

$$\mu_t = Z^m f_t, \quad (13)$$

where the $6 \times M$ matrix of factor loadings Z^m relates f_t with μ_t . As the macroeconomic variables have been standardized before they enter the analysis, we do not include a constant term in (13). The conditional score and information matrix for the Gaussian component are

$$\nabla_t^m = \left(\tilde{S}_t Z^m \right)' \left(\tilde{S}_t \Sigma_m \tilde{S}_t' \right)^{-1} \tilde{S}_t (y_t^m - \mu_t), \quad (14)$$

$$\mathcal{I}_t^m = \left(\tilde{S}_t Z^m \right)' \left(\tilde{S}_t \Sigma_m \tilde{S}_t' \right)^{-1} \tilde{S}_t Z^m. \quad (15)$$

From these quantities we can compute s_t in (3) via (6) and (7). This part of the estimation leads to a natural setting: the dynamics of the mean parameter μ_t respond linearly to the prediction error vector $y_t^m - \mu_t$ and the updating step relies on a generalized least squares computation.

Rating transition model

For the credit rating transitions, we specify a dynamic ordered logit model. Previous research on credit risk has focused on a standard multinomial specification; see, for example, Koopman, Lucas, and Monteiro (2008) and Koopman, Kraeussl, Lucas, and Monteiro (2009). The multinomial density does not take into account the fact that ratings are ordered. By using the ordered logit specification, the ordering can easily be taken into account. Also, the model then becomes a natural dynamic alternative to the static ordered probit model of CreditMetrics, which is one of the current industry standards, see Gupton and Stein (2005). However, it is far from evident how to construct an observation driven dynamic ordered logit model as opposed to a dynamic multinomial model. In particular, it is not clear what functions of the data should be chosen to drive the changes in the probabilities for the transitions. Our observation driven modeling framework solves these issues by relying on the score of the conditional log-likelihood function; here we take the score of the ordered logistic log-likelihood.

We specify the binary probability that firm i 's rating does not exceed j at the end of period t by

$$\tilde{\pi}_{ijt} = \text{P}[R_{i,t+1} \leq j] = \frac{\exp(\theta_{ijt})}{1 + \exp(\theta_{ijt})}, \quad (16)$$

where R_{it} is the rating for firm i at the start of month t , and $R_{it}, j \in \{\text{IG, BB, B, C, D}\}$. We have implicitly defined $\tilde{\pi}_{i,D,t} = 0$ and $\tilde{\pi}_{i,\text{IG},t} = 1$. From (16) it follows that the probability of a transition of firm i from rating R_{it} to $R_{i,t+1} = j$ is given by

$$\pi_{ijt} = \text{P}[R_{i,t+1} = j] = \tilde{\pi}_{ijt} - \tilde{\pi}_{i,j-1,t}. \quad (17)$$

The log-likelihood contribution at time t becomes

$$\sum_i \sum_j y_{ijt}^c \log(\pi_{ijt}), \quad (18)$$

where the first summation is over all firms, the second summation is over all five ratings, and the indicator y_{ijt}^c equals unity if firm i has moved from rating R_{it} to rating j during month t , and zero otherwise.

In our ordered logit specification, we let θ_{ijt} depend linearly on the time-varying factor f_t ,

$$\theta_{ijt} = z_{ijt}^c - Z_{it}^c f_t, \quad (19)$$

where z_{ijt}^c is a scalar and Z_{it}^c is a $M \times 1$ vector of factor loadings, which can both depend on firm specific information such as initial rating of the firm, industry sector, and time-varying financial ratios.

Based on equations (16) to (19), we obtain

$$\nabla_t^c = - \sum_i \sum_j \frac{y_{ijt}^c}{\pi_{ijt}} \cdot \dot{\pi}_{ijt} \cdot Z_{it}^c, \quad (20)$$

where

$$\dot{\pi}_{ijt} = \tilde{\pi}_{ijt}(1 - \tilde{\pi}_{ijt}) - \tilde{\pi}_{i,j-1,t}(1 - \tilde{\pi}_{i,j-1,t}), \quad (21)$$

Equation (20) is expressed in an intuitive form. The actual outcome of y_{ijt}^c is compared to its probability by considering the ratio y_{ijt}^c/π_{ijt} . In expectation, this ratio is one such that ∇_t^c has zero mean since $\sum_j \dot{\pi}_{ijt} = 0$. In the score, the ratio is weighted by the probability's sensitivity to changes in the time varying factor f_t . By summing over all ratings j , the information from all possible transitions is combined. The information matrix can be obtained by computing

$E[\nabla_t^c \nabla_t^{c'}]$. After some straightforward algebra, we obtain

$$\mathcal{I}_t^c = \sum_i n_{it} \left(\sum_j \frac{\dot{\pi}_{ij,t}^2}{\pi_{ij,t}} \right) Z_{it}^c Z_{it}^{c'}, \quad (22)$$

where n_{it} is an indicator which is unity when firm i exists at the start of period t , and zero otherwise.

As mentioned earlier, we have included time dummies in (19) for April 1982 and October 1999 in order to handle outliers due to the redefinitions of the rating categories. The redefinitions caused substantial incidental rerating activity during these months.

Loss given default model

The $K_t \times 1$ vector y_t^r of LGD's has a dimension which varies over time depending on how many firms default and on whether their LGD is recorded. As the loss given default rates y_{kt}^r for $k = 1, \dots, K_t$ are reported in percentage terms, it is common to model these with a beta distribution, such that the log-likelihood contribution at time t is given by

$$\sum_{k=1}^{K_t} (a_{kt} - 1) \log(y_{kt}^r) + (b_{kt} - 1) \log(1 - y_{kt}^r) - \log[B(a_{kt}, b_{kt})], \quad (23)$$

where a_{kt} and b_{kt} are positive scalars, and $B(a_{kt}, b_{kt})$ is the Beta function $B(a_{kt}, b_{kt}) = \Gamma(a_{kt}) \Gamma(b_{kt}) / \Gamma(a_{kt} + b_{kt})$. The mean of our k th beta distribution is denoted by μ_{kt}^r and we let it vary over time as a function of the factor f_t . To ensure that the mean remains within the $[0, 1]$ interval, we model it as

$$\log(\mu_{kt}^r / (1 - \mu_{kt}^r)) = z_k^r + Z^r f_t, \quad (24)$$

where Z^r is the $1 \times M$ vector of factor loadings and z_k^r is a scalar intercept term that represents the unconditional level of LGDs. The factor loadings are common to all defaults while the LGD level relies on the particular k th default at time t . When the data is sufficiently informative, the factor loadings can also vary with particular characteristics of the k th default.

Let $\beta_r > 0$ be an unknown scalar and define the variance of the beta distribution as

$$\sigma_{kt}^{r2} = \mu_{kt}^r \cdot (1 - \mu_{kt}^r) / (1 + \beta_r). \quad (25)$$

This specification insists that as long as the conditional mean remains within the boundaries of the unit interval, the variance remains positive. From the specifications of the mean μ_{kt} in (24) and the variance σ_{kt}^{r2} in (25), it follows that we can specify the shape parameters a_{kt} and b_{kt} as

$$a_{kt} = \beta_r \cdot \mu_{kt}^r, \quad b_{kt} = \beta_r \cdot (1 - \mu_{kt}^r)$$

The conditional score and information matrix for the beta distribution function, under our parameterization, are given by

$$\nabla_t^r = \beta_r \mu_t^r (1 - \mu_t^r) (Z^r)' (1, -1) \sum_{k=1}^{K_t} \left((\log(y_{kt}^r), \log(1 - y_{kt}^r))' - \dot{B}(a_{kt}, b_{kt}) \right) \quad (26)$$

$$\mathcal{I}_t^r = (\beta_r \mu_t^r (1 - \mu_t^r))^2 (Z^r)' (1, -1) \left(\sum_{k=1}^{K_t} \ddot{B}(a_{kt}, b_{kt}) \right) (1, -1)' Z^r \quad (27)$$

where $\dot{B}(a_{kt}, b_{kt})$ and $\ddot{B}(a_{kt}, b_{kt})$ are the first and second derivatives of the log Beta function with respect to $(a_{kt}, b_{kt})'$, respectively. The quantities ∇_t^r and \mathcal{I}_t^r enable the computation of the scaled score function that we need for our dynamic modeling framework. This observation driven time-varying parameter model for the beta distribution is new to the literature.

In our modeling framework, the three different dynamic models for the normal, the ordered probit and the beta densities are integrated naturally into an observation driven dynamic factor model. The log-likelihood value for the observations at time t is obtained by simply adding the components in (12), (18), and (23). Similarly, the score and information matrix are obtained simply by $\nabla_t^m + \nabla_t^c + \nabla_t^r$ and $\mathcal{I}_t^m + \mathcal{I}_t^c + \mathcal{I}_t^r$, respectively.

4 Empirical results

The presentation of the empirical results is split into two parts. First, we discuss the in-sample results of the model. Then, we use the model to forecast economic and credit risk scenarios

out-of-sample.

4.1 In-sample results

In our study we have considered several models with different combinations of factors. All factors are subject to the dynamic specification (3) with $p = q = 1$. Preliminary estimation results reveal that we need at least two macro factors to capture the salient features of the data. After further analyses we have come to the conclusion to consider models with two or three macro factors, one or two frailty factors, and possibly a separate LGD factor. Frailty factors are designed to capture default clustering above and beyond what is implied by shared exposure to common macroeconomic risk factors. Such excess default clustering is apparent from studies of Das, Duffie, Kapadia, and Saita (2007), Duffie, Eckner, Horel, and Saita (2009), Azizpour, Giesecke, and Schwenkler (2010) and Koopman, Lucas, and Schwaab (2011).

We impose a recursive structure on the factors for identification. In particular, we allow the macro factors to influence the macro series, the transition probabilities, and the LGD mean and variance. The frailty (or transition) factors influence only the transition probabilities and the LGD mean and variance. The frailty factors thus capture variation in default and rating transition probabilities above and beyond what is captured by macroeconomic variables alone. Finally, the LGD factor influences only the LGD dynamics and is put into the model to determine whether LGD dynamics coincide with frailty dynamics, or not. We label each model by ordering the number of factors as (m, c, r) such that a model labeled $(3, 2, 1)$ denotes a model with 3 macro factors, 2 credit risk (frailty) factors, and 1 LGD factor.

The effects on the factor loadings are straightforward. Starting with the macroeconomic variables, we restrict Z^m such that only the macro factors directly load into the conditional mean μ_t of y_t^m . The columns of Z^m associated with the frailty and LGD factors are zero. In the matrix Z^c , we allow the macro factors to enter the log-odds (19) while the columns associated with any LGD factors are set equal to zero. We leave Z^r unrestricted meaning that both macro factors and credit risk factors load into the conditional mean of the LGDs y_{kt} .

Further restrictions have to be imposed on the factor loadings Z^i for $i = m, c, r$ to ensure identification of both Z^i and the parameters $A = A_1$ and $B = B_1$ in (3). We restrict the rows in Z^m associated with industrial production, the credit spread, and realized volatility. In models

with two credit factors, we restrict the rows in Z^c associated with $R_{it} = \text{BB}$ and $R_{it} = \text{CCC}$. Furthermore, the matrices A and B are taken to be diagonal. Under this specification, the model is identified.

Although some factors do not contemporaneously load into all of the observation densities, information in series with zero loading coefficients can still affect the factors in future periods through the score function in recursive process (3). For example, the credit frailty factors and the LGD factor do not enter the equation for the conditional mean of the macros μ_t in (13), but the macro factors do load into the transformed mean and variance of the LGDs. Therefore, information in the credit ratings transition frequencies and LGDs helps to determine the value of the macro factors at time $t + 1$ because they are part of the score vector.

Table 1 contains a list of estimated models with their log-likelihood, Akaike information criterion (AIC), and Schwarz Bayesian information criterion (BIC) values. We see a marked improvement in the log-likelihood as more factors are added. Particularly adding a third macro factor or a first frailty factor causes a substantial likelihood improvement. Adding a second frailty factor further increases the likelihood, but more modestly. Adding a separate LGD factor appears to have a negligible effect on the likelihood if the LGDs already load on the macro and frailty factors.

Looking at the model selection criteria, we see that the (3,2,0) is preferred by both the AIC and BIC criterion, respectively. We therefore take the (3,2,0) model as our benchmark in the remaining analyses. Table 2 presents the estimated parameters and standard errors for this model. Standard errors are computed using the inverse Hessian of the maximized log likelihood.

The coefficients for the dynamics of the factor are strongly significant as we observed from the significance of the A coefficients in Table 2. Moreover, the macro and frailty factors are highly persistent: all of the B coefficients are estimated as 0.9 or higher. This implies that rating transition probabilities, including default probabilities, are very sticky and may deviate from their unconditional values for a substantial number of months. From a risk management perspective, this means that capital levels must be set in accordance with an episode (rather than an incidence) of high default rates for any portfolio of credit exposures.

The estimation results for Z^m reveal that the first macro factor loads on industrial production growth, real GDP growth, the (negative) change in the unemployment rate, and to a lesser extent on the (negative) credit spread. The first three variables enter with roughly

equal weights. The first macro factor can therefore be interpreted as a standard business cycle indicator. The factor is higher if the business cycle is in a good state. The results for Z^c show that the first macro factor mainly feeds significantly into the rating transition probabilities for the lower grades: bad business cycle conditions imply higher default probabilities. This holds most strongly for the CCC class, followed by the B and BB classes. The investment grade (IG) class does not load significantly on the first macro factor.

The second macro factor mainly loads on the credit spread and the equity market volatility, and to a lesser extent on the change in the unemployment rate and the negative annual stock return. All coefficients have consistent signs: high credit spreads, high volatilities, bad returns and upward changes in unemployment all push the second macro factor up. As credit spreads and volatilities are the main ingredients of the second macro factor, we interpret it as a summary of the perceptions of financial markets on economic conditions. The corresponding estimates of the elements of Z^c show that the second macro factor significantly feeds into all transition probabilities, but particularly into the transition probabilities for higher grade firms. The signs are all intuitive: if markets perceive credit risk to be high and the environment to be uncertain, we tend to see more defaults and downgrades. The second macro factor also significantly influences the mean LGD rate: high credit spreads and high volatility are positively correlated with high LGD rates.

The third macro factor is somewhat harder to interpret and loads on volatility, real GDP growth, and realized returns, and to a lesser extent on the (negative) change in the unemployment rate. The last four variables indicate that the third macro variable might be a proxy for business cycle conditions. However, its relation with realized volatility implies that higher uncertainty also pushes the factor up. The corresponding Z^c coefficient estimates reveal that they all have the same sign and again are predominantly significant for the high grade firms. A high value of the factor pushes default and downgrade probabilities down. The impact on the LGD rates is negative, but only significant at the 10% level. All in all this favors the interpretation of the third macro factor as another proxy for the business cycle climate.

We now turn to the frailty factors and their loadings in Z^c and Z^r . These factors explain autonomous default dynamics beyond what is implied by shared exposure to the common macro factors. The first frailty factor is the most important one and captures the default dynamics that are left after controlling for the macro factors. The coefficients for all three rating categories

are substantial and significant. If the frailty factor is high, default and downgrade probabilities go up. Interestingly, the frailty factor also feeds significantly into the LGD equation. Risk of credit ratings migrations and LGD risk appear to move together.

The second frailty factor mainly loads positively on the CCC firms and negatively on IG firms. The factor captures two historical features of the corporate bond market. First, in the mid 80s and also in recent years, the number of defaults of investment grade firms is higher than expected. This causes the significantly negative coefficient on the investment grade loading in Z^c . Second, via the loading of the CCC class, the factor also picks up the benign default climate in the years preceding the financial crisis. Due to its direct link with these historical default periods, the second frailty component also impacts the LGD dynamics significantly.

Based on all estimation results, we conclude that conditioning on macro factors alone is not sufficient to capture credit risk dynamics. Both transition probabilities and LGD dynamics are affected by more than only macro factors.

The final coefficients in Table 2 are the intercepts, measurement volatilities, and dummy coefficients. The intercepts z^c of the ordered logit specification show that ratings are highly persistent on a monthly basis. For example, by considering only the cut-off points of the ordered logit specification, we find that the probability of remaining in investment grade is $(1 + \exp(-6.291))^{-1} \approx 99.82\%$, while the probability of a CCC company defaulting over the next month is $(1 + \exp(3.817))^{-1} \approx 2.15\%$. As expected, most of the dummy coefficients are highly significant, which is a direct consequence of the abnormally high re-rating activities in these months due to institutional features.

Graphs of the estimated five factors are presented in Figure 1, and their implied fits for the macro series in Figure 2. We confirm our earlier finding on the relation of the first and third macro factors to growth, and of the second macro factor to credit spreads and volatilities. The fit to the macro factors is tight. Only the highest peaks in the volatility series are missed by our modeling framework. The close fit is most likely due to the dynamic factor specification of our model in combination with the monthly updating frequency of the factors.

The transition probabilities are presented in Figure 3. The probabilities are driven by both the macro and frailty factors. The first frailty factor in Figure 1 captures most of the dynamics in defaults and downgrades. In particular, we see the high peak in defaults around 1990-1991 and 2000-2002. Interestingly, the first frailty factor is low in the most recent years of the sample,

implying that the number of defaults is lower than what should be expected based on macro fundamentals.

As mentioned before, the second frailty factor in Figure 1 captures higher downgrade and default rates during the savings and loans and the financial crisis, as well as lower default rates of CCC graded companies in the wake of the dotcom crisis. The impact of the combined macro and frailty factors on transition probabilities in Figure 3 shows the structure of the ordered logit model. In each row, the dynamics of ratings are driven by the same linear combination of unobserved factors. Given that we consider transitions at the monthly frequency, all probabilities are close to zero, except the probability that the rating remains unchanged. The peaks for the different ratings differ across rating grades. For example, the investment grade class has its highest default peaks in the financial and the dotcom crises, and some substantially lower peaks in the mid 80s and the 1991 recession. By contrast, the CCC class has its highest peak in 1991. However, even though the magnitude of the peaks and troughs differs across rating grades, all rating classes appear to share the common dynamics of clustered default experience present in the data.

Finally, we turn to the LGD series. In the upper-right panel of Figure 4 it appears that the LGD data is very noisy and irregularly spaced. During some months, we do not observe any LGDs, while in other months we observe many defaults and LGDs. The extracted factors from our model appear to represent the salient features of our data set well. LGDs correlate positively with default rate dynamics via the common factors in the model. Indeed, defaults and LGDs tend to be high and low at the same time.

The effect of the time variation in the parameters of the beta distribution is clearly visualized in the left hand panels of Figure 4. The upper-left panel gives the cross-sectional beta distributions for LGDs applicable in a mild year, June 2006. We first concentrate on the (3,2,0) model, indicated by the solid-bar histogram. In June 2006, the bulk of the probability mass for the (3,2,0) model is to the left, indicating that LGDs are typically low at that time. By contrast, the result for January 2009 in the lower-left panel are quite the opposite. Most of the probability mass has shifted to the right, indicating that LGD rates have gone up substantially on average.

We conclude our discussion of the in-sample results by focusing on the difference between a model with (3,2,0) and without (3,0,0) frailty effects. The upper-right graph indicates that

the effect on the dynamics of LGDs is substantial. Only accounting for macro effects and discarding potential frailty, the model's fit misses both historically high and low LGD rates. This is corroborated in the left-hand graphs. During June 2006, the implied beta distribution for LGDs for the (3,0,0) is substantially different than the one for the (3,2,0) model. During other periods such as January 2009, the differences are small.

The different implications of adding frailty for credit risk are also evident if we consider transition probabilities. The lower-right panel in Figure 4 presents the transition probability from BB to default. We find that the dynamics are substantially different between the two models, and particularly that the model with only macro factors misses some of the salient peaks and troughs present in the (3,2,0) model. For example, the (3,0,0) model misses the fact that the default rates in the years leading up to the financial crises are atypically low compared to macro fundamentals.

4.2 Forecasting and credit risk management

In this section, we use our observation driven mixed measurement dynamic factor model to simultaneously forecast the dynamics of macro variables, rating transition and default probabilities, and LGDs for credit portfolio loss distributions. Such forecasting distributions for portfolio losses are an important ingredient for assessing and managing the risk of credit exposures; see Lando (2004). We can carry this out by, for example, having internal economic capital requirements based on portfolio credit loss distributions.

We imagine a scenario where at time T , the end of the sample, a financial institution holds a large portfolio of bonds. The risk management department would like to know the distribution of credit losses at a specified future time $T+H$ to set up capital requirements or hedge positions. To simplify the simulation experiment, we concentrate on default losses only and do not include credit losses due to downgrades in a marked-to-market setting. Such an extension is possible if one wants to assume a credit spread model for different maturities and rating grades. The set-up is close to treating the portfolio as a banking book in the accounting and regulatory sense. We also assume a zero discounting rate to cumulate losses that materialize at different points in time. The model can be extended by including a full term structure component. Modeling the full term structure is, however, not the key focus of our study.

As in the previous subsection, the initial portfolio of bonds was set equal to the current exposures in the last period of the observed data. This leads to a portfolio of 1144 firms rated IG, 265 firms rated BB, 615 firms rated B, and 311 firms rated CCC.

Given the portfolio and our parameter estimates, our modeling framework can be used to generate the dynamic evolution of the macro variables and the rating composition of the portfolio via simulation. If one of the firms transits into default, the LGD component in our model can be used to generate the corresponding LGD realization of the appropriate beta distribution. By simulating the portfolio and macros forward in this way, we obtain a realization of default losses between T and the horizon date $T + H$. This process can be repeated many times. In our set-up below, we use 500,000 simulations for the loss distribution at different horizons.

The model specification and starting values f_T are an important ingredient of the analysis. In particular, we compare a model with and without frailty dynamics. In addition, we consider starting values in a benign period and in a period of stress. Given the stickyness of the macro and frailty factors, these different starting conditions have a different impact on the loss experience, both at short and longer horizons.

Figure 5 presents our initial results. It reveals the simulated loss distributions at four different horizons for the (3,2,0) model. At the one month horizon (upper-left), we clearly visualize the differences in the forecasting distribution. Starting from a recession, next month's losses are higher on average and are also more spread out. If current economic conditions are good, by contrast, next period's losses are low on average and have a smaller standard deviation. Even at the 1-month horizon, there is a significant non-overlap of the 99% confidence regions of the two densities. Both the macros and the frailty factors are in a state of recession versus expansion at time T , causing the PDs and average LGDs of the next month to be substantially different.

As the forecasting horizon H increases, the densities gradually start to overlap more and more. The effect is due to the stationarity of the model and the high persistence of the factors in f_t . The high persistence of the unobserved components f_t causes the PDs and expected LGDs to be substantially different for a number of consecutive months. The cumulated losses over longer time periods remain substantially different. Over longer and longer horizons, however, the stationarity of the model (mean reversion) takes over and the influence of initial conditions

starts to vanish. This can be viewed at a horizon of three years. Over a 36 months horizon, a current recession might easily turn into an expansionary phase, and vice versa.

We compare the economic contribution of the frailty factors to capital requirements in Figure 6. The same approach is followed as for Figure 5, but for the (3,2,0) and the (3,0,0) model, respectively. For the left-hand four panels in the figure, the factors are started at their unconditional means, $f_t = 0$. The panels show that the forecasting distributions of models with and without the two frailty factors are roughly similar at the one-month horizon. If anything, the probability of large losses is somewhat larger for the model with frailty factors (3,2,0).

As we start to focus on longer horizons, the differences become much clearer. The models that only use the three macro factors to drive the credit loss conditions result in much more concentrated loss distributions. The location of these distributions is roughly the same as that of the model with the frailty components, but the spread and right hand tail are substantially different. This is most clearly seen at the 36-month horizon: the (3,2,0) model has the right skew that is typical for portfolio loss distributions. Because of the additional, separate dynamics of the frailty components versus the macro components, the 99% confidence loss quantiles for the (3,2,0) are substantially larger than for the (3,0,0) model. A model that only conditions on macro risk leads to a severe underestimation of required capital, particularly at longer horizons.

The four right-hand panels in Figure 6 give the results if the factors f_T are started in a recession. At short horizons, the results are now much more pronounced. This is due to the fact that not only the macro factors start under bad economic circumstances but the frailty factors do as well. The projected increase in expected loss rates are substantial. The upper quantiles of the loss distribution increase substantially if the frailty factors are added. The increase for the highest quantiles is close to 100% in most settings, implying a doubling of model based capital requirements at these horizons. At the longest horizon of 36 months, the distributions start to overlap more, similar to the setting with $f_T = 0$. This is again due to the stationarity of the model: at these longer horizons, the differences in initial conditions matter less. Still, the loss quantiles are substantially different for models with and without frailty.

We conclude that models that account for only macro dynamics can miss out on substantial parts of credit loss dynamics and potential credit loss sizes. Though the absolute numbers for the losses (as a percentage of the notional) may appear small, one should account for the fact that the portfolio holds many investment grade bonds, resulting in a high quality portfolio with

small losses on average. If we look at the 99% quantile of the loss distribution at long horizons, accounting for the frailty dynamics may shift out the loss quantiles by more than 80%.

4.3 Impulse response analysis

Another application of our observation driven mixed measurement model is impulse response analysis that allows us to track how a shock in an unobserved factor process feeds through the entire model. In particular we are interested in how such shocks dynamically affect the credit loss distributions.

Since our dynamic factor model is non-linear and non-Gaussian, we follow the non-linear impulse response methodology of Koop, Pesaran, and Potter (1996). The approach is as follows. Consider the (3,2,0) model where we keep the parameter estimates fixed at their maximum likelihood values. We store the estimated scaled scores $\hat{s}_1, \dots, \hat{s}_T$. Our aim is to compute

$$\mathbb{E}[g(y_{T+H})|s_T = s^*, \mathcal{F}_{T-1}] - \mathbb{E}[g(y_{T+H})|\mathcal{F}_{T-1}], \quad (28)$$

for some function $g(\cdot)$. Concretely, this means we compare the effect on $g(y_{T+H})$ of shocking the common factor f_T by a fixed shock s^* versus the average effect over all possible shock sizes. To compute the second expectation in (28), we bootstrap a value s_{T+1} from the fitted scaled scores $\hat{s}_1, \dots, \hat{s}_T$. Together with f_T , this draw is used to generate f_{T+1} , which in turn can be used to simulate the complete model forward (macros, rating transitions of all firms in the entire portfolio, and LGDs in case of defaults). We store the simulated portfolio losses for each of the next 48 months. The first expectation in (28) is obtained similarly. Suppose we are interested in a shock of size s^* to the i -th factor within f_t . After drawing a bootstrapped value of \hat{s}_t , the i -th element is set equal to s^* , which in our case is a one unit negative shock. The impulse responses obtained in this way have marginalized out all other elements besides the i -th entry. In our work, we use 50,000 simulations. Plots of the impulse responses of the loss distribution are presented in Figure 7 and 8.

The impulse response functions for the macros confirm the estimation results as presented in Table 2.¹ The impulse response function for the first macro factor (as presented in the first

¹We have smoothed the plots because even with a high number of simulations the impulse responses exhibit jagged behavior due to the high quality of the portfolio (few defaults) and the discrete nature of defaults.

column of figures) has the largest impact on the business cycle related variables. A bad shock decreases industrial production growth (1st row) and real GDP growth (3rd row) and increases the unemployment rate (2nd row). The impact is quite persistent and roughly needs 3 to 4 years to die out. The second factor mainly impacts the lower three macros: the credit spread, the stock market return, and the volatility. Again, the effect of a shock in the second factor dies out only slowly: its impact lasts 3 to 4 years. The third factor appears to be of a much more transitory nature and impacts the stock market return and its volatility (and to a lesser extent real GDP growth).

As expected, the frailty factors have no direct impact on the macro variables. This underlines the recursive structure of the model and helps to interpret the frailty factors as credit dynamics above and beyond macro developments.

The results for the credit loss distributions and their 90% quantiles are in Figure 8. Both the means and quantiles, however, reveal the same pattern. We see that the first and second macro factor have the largest and most persistent impact on portfolio credit losses. The effect of a shock in the first or second macro factor can last up to three or four years. The third macro factor has a much shorter half-life and vanishes after one to two years. Its impact on portfolio credit losses is also substantially smaller.

The frailty factors have a large impact on credit losses. The first frailty factor (column 4) has an effect similar in size to the business cycle related macro factor (column 1) and the financial markets related macro factor (column 2). The impact of the second frailty factor is roughly half this size. This implies that by omitting the frailty component from the model, one may miss one third to one half of the credit risk. In addition, the dynamics of credit losses might be completely misspecified if the frailty components are left out of the model.

We conclude that the in-sample estimation results, the out-of-sample forecasting distributions, as well as the impulse response analysis all point in the same direction: the size and dynamics of portfolio credit losses cannot be captured by conditioning only on macroeconomic conditions. The losses appear to have their own additional (frailty) dynamics, which are highly significant in both statistical and economic terms. The current observation driven mixed measurements modeling framework allows for a straightforward analysis of all these results. It does so in a standard likelihood context, where the likelihood is easily tractable and known in closed form. As such, the framework provides a good alternative to parameter driven models that

require advanced simulation tools to do estimation and inference.

5 Conclusion

We have introduced a new framework for observation driven mixed-measurement dynamic factor models for time series observations from different families of distributions and mixed sampling frequencies. Missing values arise due to unbalanced time series panels and mixed frequencies; they can be accommodated straightforwardly in our framework. In an empirical application of the mixed-measurement framework we model the systematic variation in US corporate default counts and recovery rates in the period 1982–2010. We estimate and forecast interconnected default and recovery risk conditions, and demonstrate how to obtain the predictive credit portfolio loss distribution.

The model is a useful device to deal with a large number of data points in a complex data set. A clear advantage of our framework is that the likelihood remains analytically tractable in closed form and therefore standard likelihood procedures can be used for parameter estimation. The model also lends itself easily to integrated forecasting exercises for joint macro and credit risk developments. In particular, the model can be used in a straightforward way to obtain portfolio loss distributions at multiple horizons. Such distributions can be used as an input for risk analyses. In addition, we can use the simulation framework to conduct impulse response analysis for non-linear model specifications. The impulse response functions can be used directly to study the feed-through mechanism from macro developments to credit losses and thus provide interesting input for both practitioners, regulators, and policy makers.

We have shown that the mixed-measurement dynamic factor modeling framework facilitates our credit risk application well given the many complexities in a typical credit risk data set. However, the modeling framework is not restricted to credit risk applications. For example, in the context of high-frequency financial data, our approach allows to model inter-trade durations, discrete tick changes in prices, and general market conditions simultaneously for different assets that can be made subject to common risk and liquidity factors. For the modeling of a macroeconomic time series panel, we can mix the usual continuous variables with indicator data such as the NBER business cycle classifications. Other examples may be found in health and retirement economics, business, marketing, and psychology.

References

- Allen, L. and A. Saunders (2004). Incorporating systematic influences into risk measurements: a survey of the literature. *Journal of Financial Services Research* 26(2), 161–191.
- Altman, E. I., B. Brady, A. Resti, and A. Sironi (2003). The link between default and recovery rates: Theory, empirical evidence, and implications. *Journal of Business* 78, 2203–2227.
- Azizpour, S., K. Giesecke, and G. Schwenkler (2010). Exploring the sources of default clustering. *Stanford University working paper series*.
- Bai, J. (2003). Inferential Theory for Factor Models of Large Dimension. *Econometrica* 71(1), 135–171.
- Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica* 70(1), 191–221.
- Bai, J. and S. Ng (2007). Determining the number of primitive shocks in factor models. *Journal of Business and Economic Statistics* 25(1), 52–60.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31, 307–327.
- Connor, G. and R. A. Korajczyk (1988). Risk and Return in equilibrium APT: Application of a new test methodology. *Journal of Financial Economics* 21(2), 255–289.
- Connor, G. and R. A. Korajczyk (1993). A test for the number of factors in an approximate factor model. *Journal of Finance* 48(4), 1263–91.
- Cox, D. R. (1981). Statistical analysis of time series: some recent developments. *Scandinavian Journal of Statistics* 8, 93–115.
- Creal, D. D., S. J. Koopman, and A. Lucas (2010). Generalized Autoregressive Score Models with Applications. Working paper, University of Chicago Booth School of Business.
- CreditMetrics (2007). CreditMetrics (TM) - Technical Document, RiskMetrics Group. www.riskmetrics.com/pdf/dnldtechdoc/CMTD1.pdf.
- Das, S. R., D. Duffie, N. Kapadia, and L. Saita (2007). Common failings: how corporate defaults are correlated. *The Journal of Finance* 62(1), 93–117.
- Doz, C., D. Giannone, and L. Reichlin (2006). A Quasi Maximum Likelihood Approach for Large Approximate Dynamic Factor Models. *ECB working paper* 674.
- Duffie, D., A. Eckner, G. Horel, and L. Saita (2009). Frailty correlated default. *The Journal of Finance* 64(5), 2089–2123.
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50(4), 987–1007.
- Engle, R. F. (2002). Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics* 20(3), 339–350.

- Engle, R. F. and G. M. Gallo (2006). A multiple indicators model for volatility using intra-daily data. *Journal of Econometrics* 131, 3–27.
- Engle, R. F. and J. R. Russell (1998). Autoregressive conditional duration: a new model for irregularly spaced transaction data. *Econometrica* 66(5), 1127–1162.
- Forni, M., M. Hallin, M. Lippi, and L. Reichlin (2000). The Generalized Dynamic-Factor Model: Identification and Estimation. *The Review of Economics and Statistics* 82(4), 540–554.
- Forni, M., M. Hallin, M. Lippi, and L. Reichlin (2005). The Generalized Dynamic-Factor Model: One-Sided Estimation and Forecasting. *Journal of the American Statistical Association* 100, 830–840.
- Geweke, J. (1977). The dynamic factor analysis of economic time series. In D. Aigner and A. Goldberger (Eds.), *Latent Variables in Socio-Economic Models*, pp. 365–383. Amsterdam: North-Holland.
- Ghysels, E., P. Santa Clara, and R. Valkanov (2006). Predicting volatility: getting the most out of return data sampled at different frequencies. *Journal of Econometrics* 131(1-2), 59–95.
- Gupton, G. and R. Stein (2005). LossCalc V2: Dynamic Prediction of LGD. *Moody's KMV Investors Services*.
- Jungbacker, B. and S. J. Koopman (2008). Likelihood-based Analysis of Dynamic Factor Models. *Tinbergen Institute Discussion Paper*.
- Koop, G., M. H. Pesaran, and S. M. Potter (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics* 74(1), 119–147.
- Koopman, S. J., R. Kraeussl, A. Lucas, and A. Monteiro (2009). Credit cycles and macro fundamentals. *Journal of Empirical Finance* 16(1), 42–54.
- Koopman, S. J., A. Lucas, and A. Monteiro (2008). The multi-state latent factor intensity model for credit rating transitions. *Journal of Econometrics* 142(1), 399–424.
- Koopman, S. J., A. Lucas, and B. Schwaab (2011). Modeling frailty correlated defaults using many macroeconomic covariates. *Journal of Econometrics*, forthcoming.
- Lando, D. (2004). *Credit Risk Modeling: Theory and Applications*. Princeton, NJ: Princeton University Press.
- Mariano, R. S. and Y. Murasawa (2003). A new coincident index of business cycles based on monthly and quarterly series. *Journal of Applied Econometrics* 18(4), 427–443.
- McNeil, A., R. Frey, and P. Embrechts (2005). *Quantitative Risk Management*. Princeton, New Jersey: Princeton University Press.
- McNeil, A. and J. Wendin (2007). Bayesian inference for generalized linear mixed models of portfolio credit risk. *Journal of Empirical Finance* 14(2), 131–149.
- Merton, R. (1974). On the pricing of corporate debt: the risk structure of interest rates. *The Journal of Finance* 19(2), 449–470.
- Sargent, T. J. and C. R. Sims (1977). Business cycle modeling without pretending to have too much a priori economic theory. *Federal Reserve Bank of Minneapolis*.

- Schuerman, T. (2006). What do we know about Loss-given-default? In D. Shimko (Ed.), *Credit Risk Models and Management*. Risk Books, London, UK.
- Stock, J. and M. Watson (2002). Macroeconomic Forecasting Using Diffusion Indexes. *Journal of Business and Economic Statistics* 20, 147–162.

Table 1: *Likelihoods and Information Criteria*

The table contains the log-likelihood values and information criteria for alternative model specifications. Each model contains a different number of macroeconomic (m), credit or frailty risk (c), and LGD factors (r) which are ordered as (m, c, r) . The maximum log-likelihood value and the minimum AIC and BIC are denoted in bold.

	(2,0,0)	(2,1,0)	(2,2,0)	(3,0,0)	(3,1,0)	(3,2,0)	(3,1,1)	(3,2,1)
log-Lik	-40447.9	-40199.1	-40162.8	-40056.2	-39817.1	-39780.8	-39812.6	-39780.0
AIC	81005.9	80520.1	80457.0	80242.4	79776.2	79713.6	79771.2	79716.0
BIC	81640.0	81223.0	81218.0	80991.0	80594.0	80589.0	80612.0	80615.0

Table 2: *Parameter estimates and standard errors for the (3,2,0) model.*

This table contains the estimated parameters and their standard errors for our model with (3,2,0) factor structure. The macros are ordered from $i = 1, \dots, 6$ as industrial production growth (IP), unemployment rate change (UR), annual real GDP growth (GDP), credit spread (CrSPR), annual return on the S&P500 (SP500), and annual realized volatility of the S&P500 returns using the past 252 daily trading days ($\sigma_{S\&P}$). Significance at the 10%, 5%, and 1% level is denoted by *, **, and ***, respectively.

	macro ₁	macro ₂	macro ₃	frailty ₁	frailty ₂				
A	0.240*** (0.012)	0.303*** (0.017)	0.278*** (0.023)	0.033*** (0.005)	0.036*** (0.012)				
B	0.968*** (0.013)	0.933*** (0.021)	0.900*** (0.025)	0.978*** (0.011)	0.973*** (0.017)				
Z^m									
	macro ₁	macro ₂	macro ₃	frailty ₁	frailty ₂	z^m	Σ_m		
IP	1.000	0.000	0.000	0.000	0.000	0.000	0.063*** (0.006)		
UR	-0.892*** (0.037)	0.122*** (0.041)	-0.062* (0.040)	0.000	0.000	0.000	0.118*** (0.012)		
RGDP	0.811*** (0.066)	0.072 (0.079)	0.336*** (0.074)	0.000	0.000	0.000	0.278*** (0.034)		
Cr.Spr.	-0.169** (0.085)	1.000	0.000	0.000	0.000	0.000	0.059*** (0.005)		
$r_{S\&P}$	0.049 (0.093)	-0.268*** (0.081)	1.223*** (0.093)	0.000	0.000	0.000	0.154*** (0.012)		
$\sigma_{S\&P}$	-0.007 (0.107)	0.648*** (0.084)	1.000	0.000	0.000	0.000	0.541*** (0.045)		
Z^c									
$R_{it} =$	macro ₁	macro ₂	macro ₃	frailty ₁	frailty ₂	z^c			
						$\leq \text{IG}$	$\leq \text{BB}$	$\leq \text{B}$	$\leq \text{CCC}$
IG	-0.052 (0.059)	0.202*** (0.055)	-0.123** (0.069)	1.475*** (0.371)	-1.165** (0.555)	6.291*** (0.156)	8.279*** (0.175)	9.617*** (0.225)	10.284*** (0.276)
BB	-0.078** (0.037)	0.172*** (0.037)	-0.102*** (0.040)	1.000	0.000	-5.204*** (0.086)	4.621*** (0.082)	7.047*** (0.123)	7.444*** (0.140)
B	-0.184*** (0.035)	0.162*** (0.031)	-0.142*** (0.040)	0.970*** (0.156)	-0.016 (0.158)	-7.641*** (0.135)	-5.461*** (0.085)	4.641*** (0.080)	5.967*** (0.089)
CCC	-0.262*** (0.057)	0.073* (0.050)	-0.018 (0.075)	1.936*** (0.465)	1.000	-9.008*** (0.399)	-8.091*** (0.291)	-5.278*** (0.194)	3.817*** (0.185)
Z^r									
	macro ₁	macro ₂	macro ₃	frailty ₁	frailty ₂	z^r	β		
Z^r	0.018 (0.049)	0.276*** (0.046)	-0.082* (0.062)	1.212*** (0.376)	1.065*** (0.301)	0.194* (0.142)	2.569*** (0.094)		
DUM									
	$t = 15$		$t = 308$						
IG	3.619*** (0.161)		-1.105 (0.988)						
BB	1.359*** (0.337)		1.813*** (0.269)						
B	-3.633*** (0.526)		-0.937*** (0.281)						
CCC	-3.229*** (1.063)		-3.058*** (0.158)						

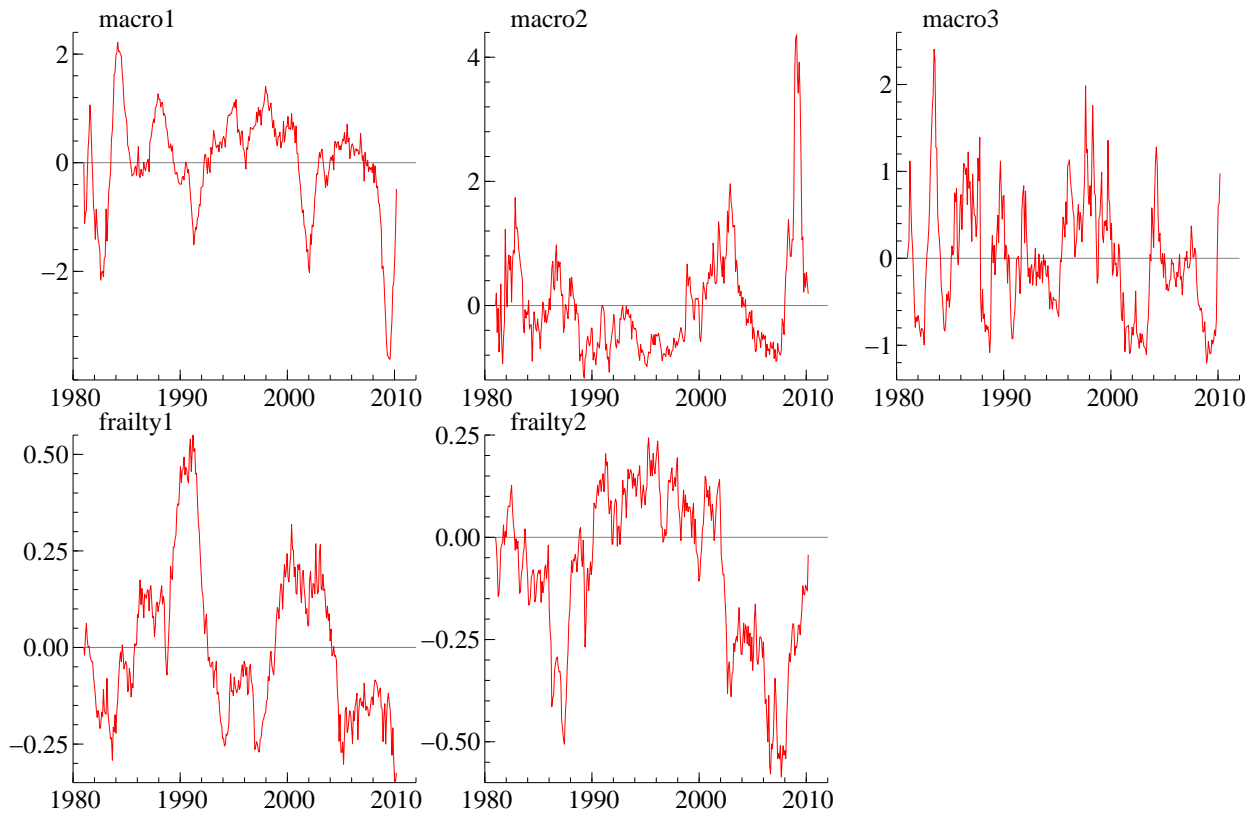


Figure 1: *Estimates of macro and frailty factors in the (3,2,0) model*

This figure contains the estimated factors f_t for a specification with 3 macro and 2 frailty factors.

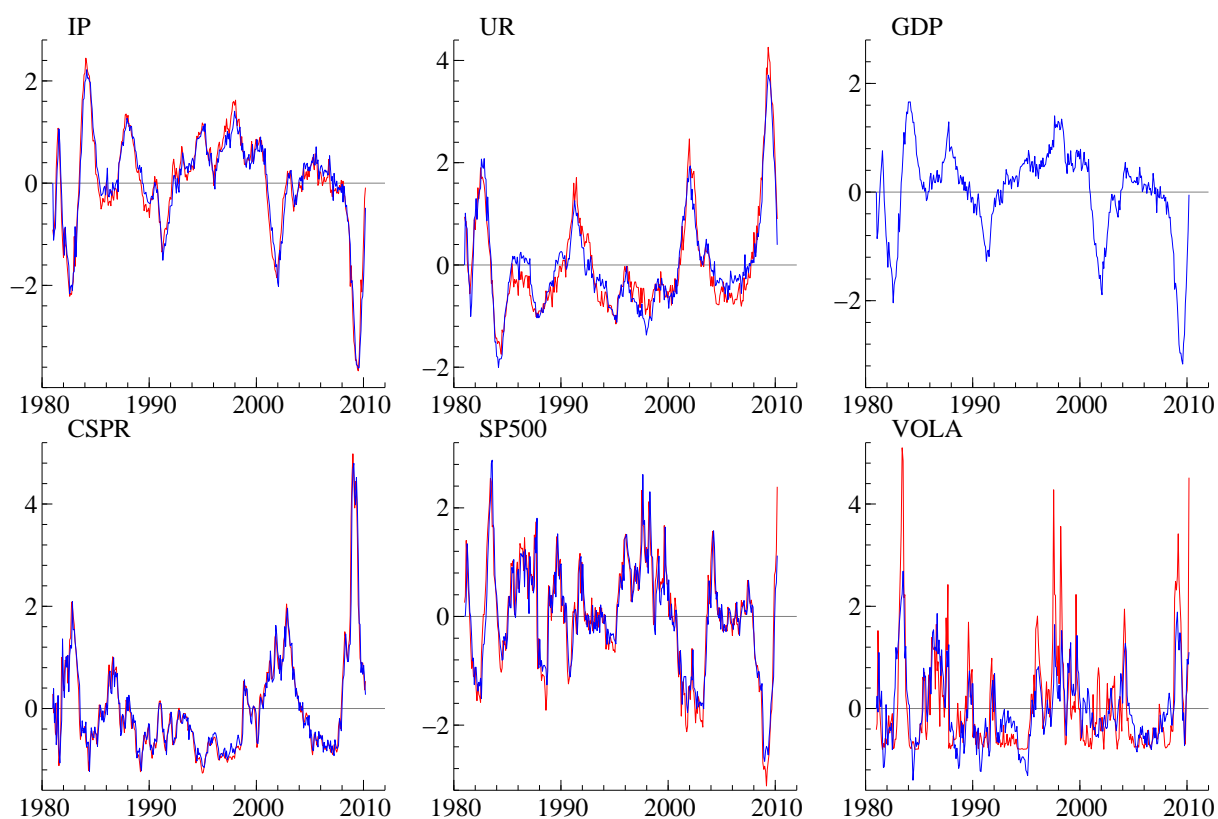


Figure 2: *Fit to macro variables*

The figure presents industrial production annual growth (IP), the annual change in the unemployment rate (UR), real GDP annual growth (GDP), the credit spread (CSPR), annual return on the S&P500 index (SP500), and realized monthly volatility (VOLA) on the S&P500 (based on daily data). Each panel contains the macro series and the fit of our model with three macro and two frailty factors, (3,2,0).

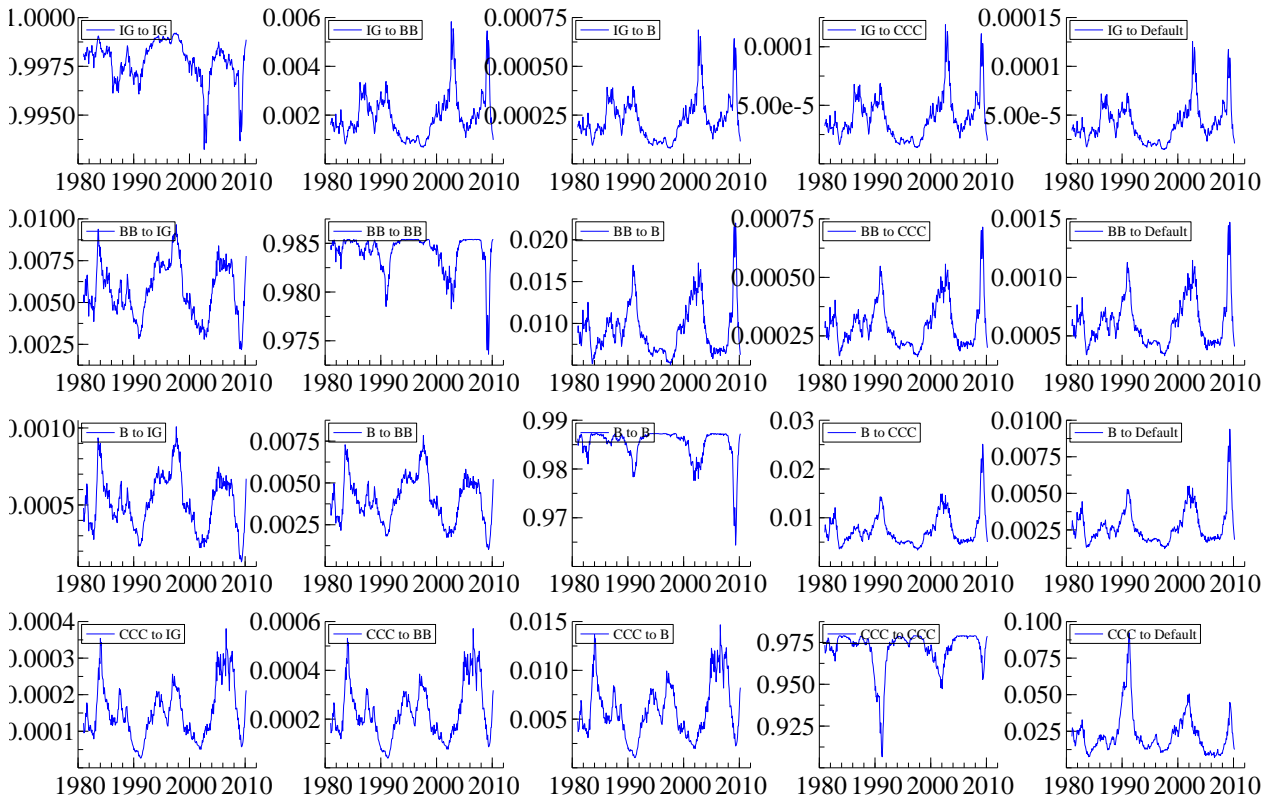


Figure 3: *Time-varying transition probabilities for the $(3,2,0)$ model*

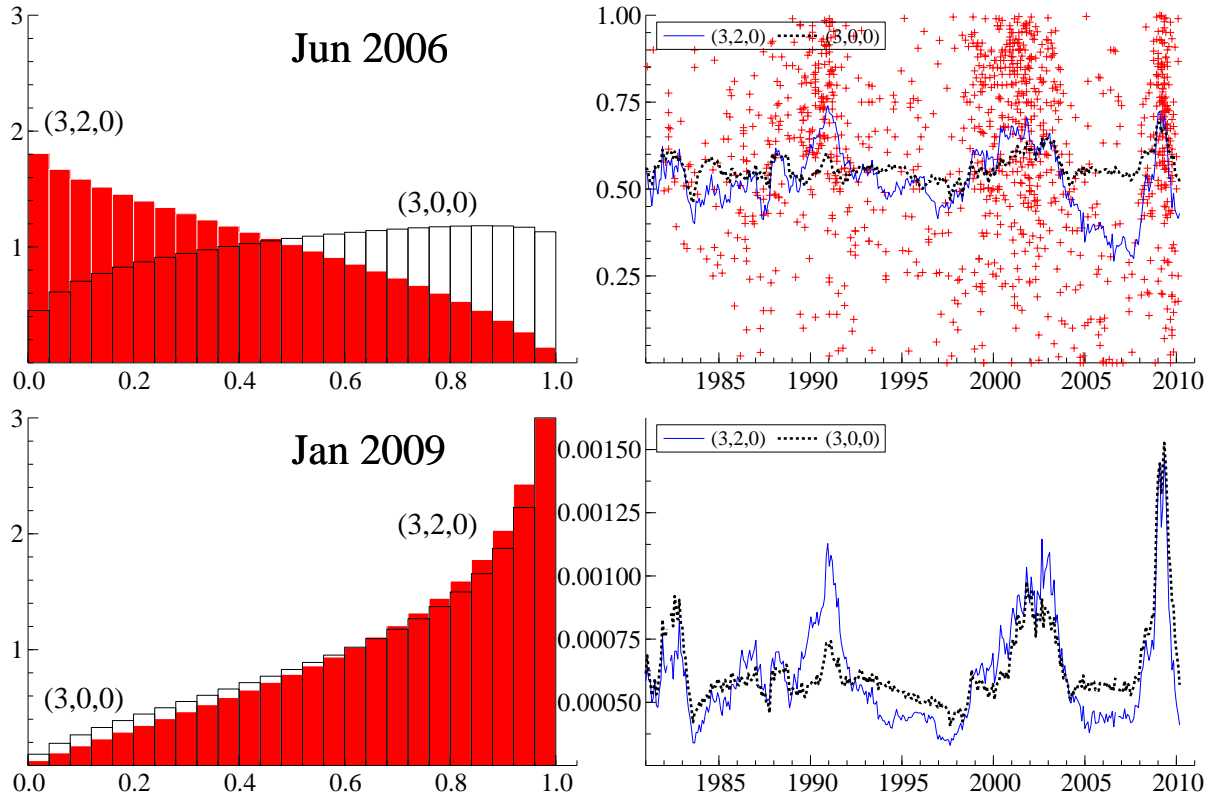


Figure 4: *Loss-Given-Default (LGD) dynamics*

The left panels contain the cross sectional beta distributions applicable in June 2006 and January 2009 for a model with three macro factors $(3,0,0)$ and a model with three macro and two frailty factors $(3,2,0)$. The upper-right panel contains the time series plot of the means of the LGD distributions and its fit to the observed LGD data. The lower-right panel gives the transition probability from BB to Default for the $(3,0,0)$ and $(3,2,0)$ models.

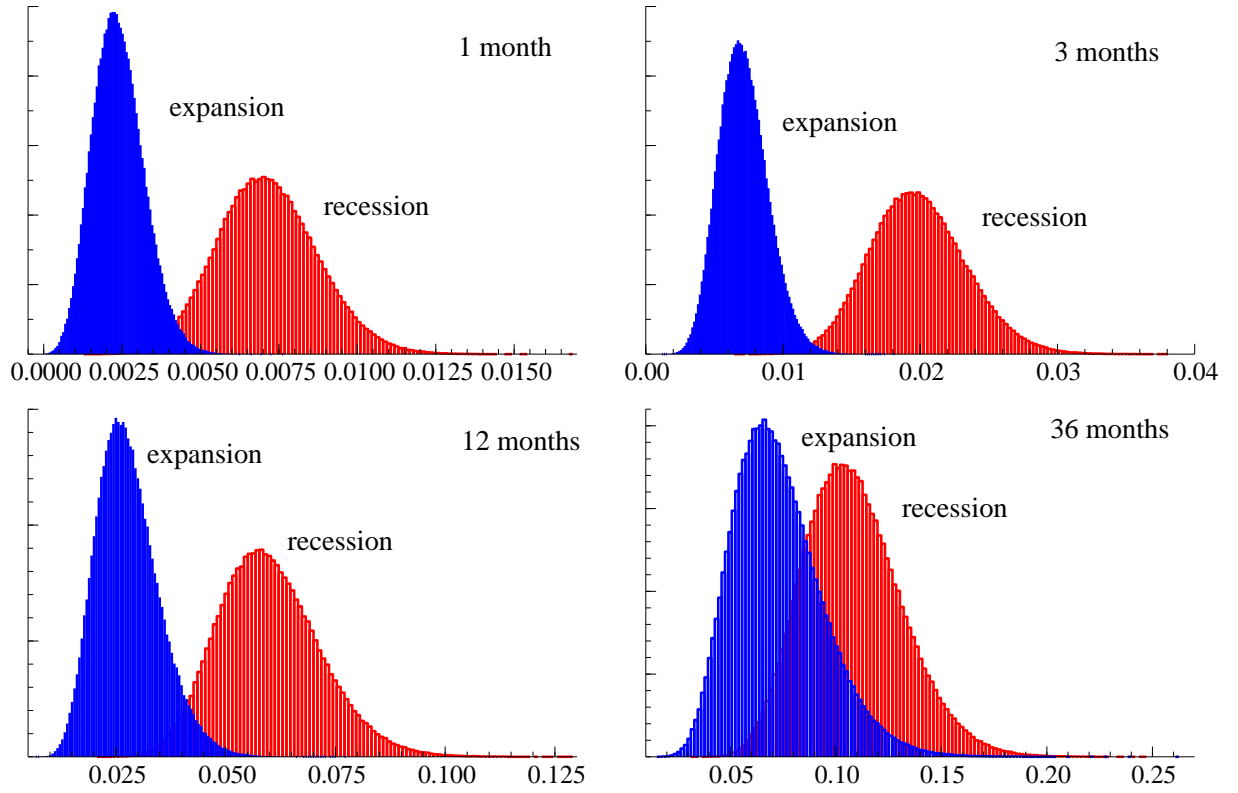


Figure 5: *Comparison of simulated loss distributions for the $(3,2,0)$ model*

For our model with three macro and two frailty factors, the panels present the cumulative losses at different horizons. The difference between the two curves is the starting values for the factors, namely recession and expansion.

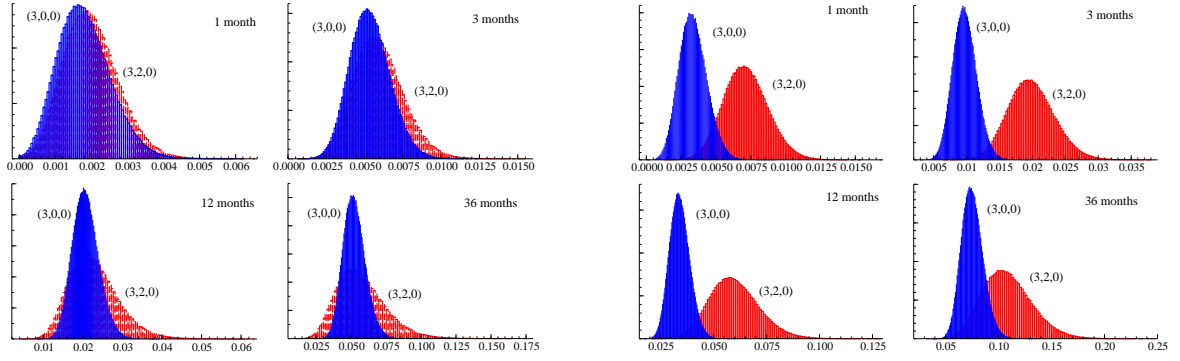


Figure 6: *Comparison of simulated loss distributions for the $(3,2,0)$ and $(3,0,0)$ model*

For our model with three macro and two frailty factors $(3,2,0)$, respectively only three macro factors $(3,0,0)$, the panels present the cumulative losses at different horizons. The left-hand four panels show the results if the factors f_t are started at zero. The right-hand four panels show the results if the factors are started in a recession period.

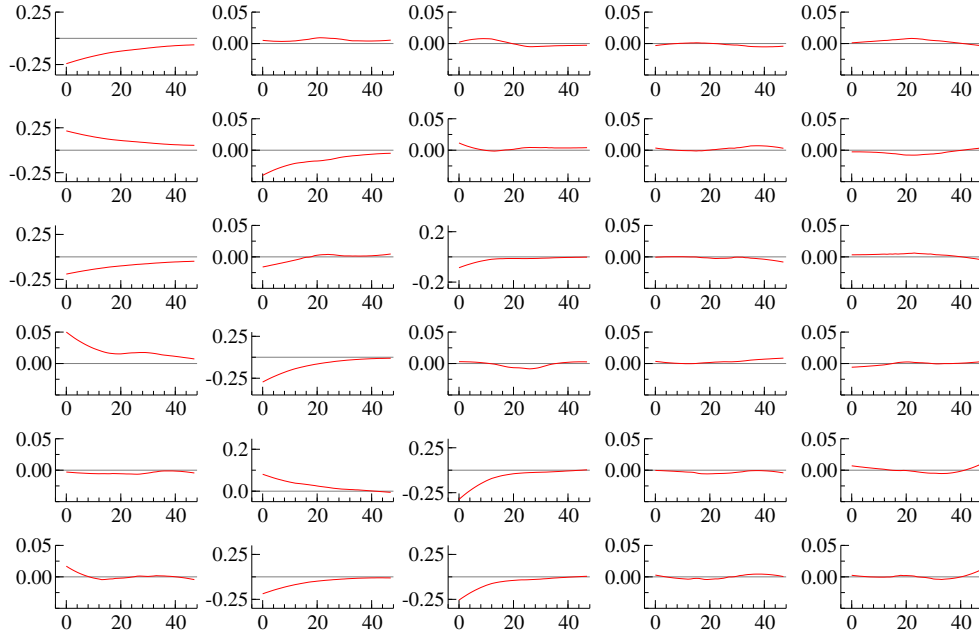


Figure 7: *Non-linear impulse response functions for the (3,2,0) model: macros*

For our model with three macro and two frailty factors, one of the factors f_t is given a unit size negative shock. All of the remaining stochastics of the model are simulated 48 months forward. The impulse response functions plot the difference between the average of the simulated quantity for a unit size shock to one of the factors and the average of the same quantity where the same factor receives a random model shock. The panels show the results for the 6 macros: rows 1 to 6 for industrial production growth, the change in the unemployment rate, real GDP growth, the credit spread, the S&P500 return, and its volatility, versus columns 1 to 5 for shocking the macro factors 1 to 3 and the 2 frailty factors. Factors are started at their mean values $f_T = 0$.

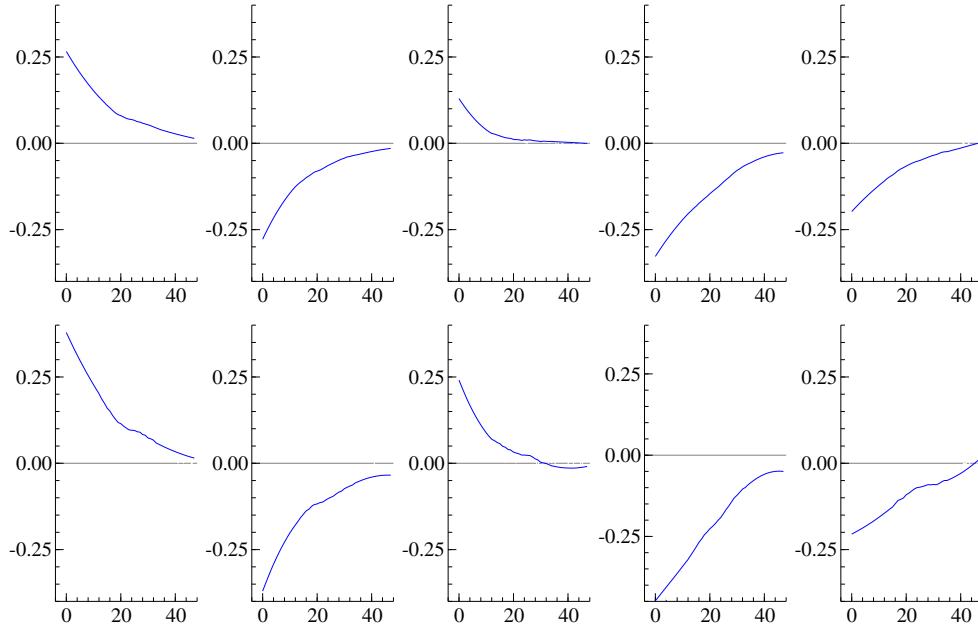


Figure 8: *Non-linear impulse response functions for the $(3,2,0)$ model: portfolio losses*

For our model with three macro and two frailty factors, one of the factors f_t is given a unit size shock. All of the remaining stochastics of the model are simulated 48 months forward. The impulse response functions plot the difference between the average of the simulated quantity for a unit size shock to one of the factors and the average of the same quantity where the same factor receives a random model shock. The panels show the result for the mean portfolio credit loss (top row) and its 90th percentile (bottom row). The portfolio holds 1144 firms rated IG, 265 firms rated BB, 615 firms rated B, and 311 firms rated CCC. Columns 1 to 5 are for a shock to the macro factors 1 to 3 and the 2 frailty factors, respectively. Factors are started at their values fitted at the end of our sample.