

Dynamic factor models with macro, frailty, and industry effects for U.S. default counts: the credit crisis of 2008*

Siem Jan Koopman^(a,c) *André Lucas*^(b,c,d) *Bernd Schwaab*^(e)

^(a) Department of Econometrics, VU University Amsterdam

^(b) Department of Finance, VU University Amsterdam

^(c) Tinbergen Institute

^(d) Duisenberg school of finance

^(e) European Central Bank, Financial Markets Research

June 1, 2011

*Author contact information: Siem Jan Koopman, VU University Amsterdam, Department of Econometrics, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, Email: s.j.koopman@feweb.vu.nl. Andre Lucas, VU University Amsterdam, Department of Finance and Financial Sector Management, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, Email: alucas@feweb.vu.nl. Bernd Schwaab, European Central Bank, Kaiserstrasse 29, 60311 Frankfurt, Germany, Email: Bernd.Schwaab@ecb.int. We thank four anonymous referees, Arthur Korteweg (our discussant at the 2011 AFA meeting in Denver), Kasper Roszbach (discussant at the Atlanta Fed), Kay Giesecke, Philipp Hartmann, David Lando, and Michel van der Wel for comments. We further thank participants at the SoFiE "Credit Risk, Liquidity, and Extreme Events" conference at the University of Chicago, the "Recent Developments in Financial Econometrics" conference at Humboldt University in Berlin, the Netherlands Econometrics Study Group meeting in Amsterdam, the "Credit Risk, Financial Crises and the Macroeconomy" 2009 CREDIT conference in Venice, the German Finance Association meeting in Frankfurt, the "Day ahead" conference at the Atlanta Fed, and seminar participants at the European Central Bank, the Swedish National Bank, Tinbergen Institute, and VU University Amsterdam. The views expressed in this paper are those of the authors and they do not necessarily reflect the views or policies of the European Central Bank or the European System of Central Banks.

Dynamic factor models with macro, frailty, and industry effects for U.S. default counts: the credit crisis of 2008

Abstract

We develop a high-dimensional and partly nonlinear non-Gaussian dynamic factor model for the decomposition of systematic default risk conditions into a set of latent components that correspond with macroeconomic/financial, default-specific (frailty), and industry-specific effects. Discrete default counts together with macroeconomic and financial variables are modeled simultaneously in this framework. In our empirical study based on defaults of U.S. firms, we find that approximately 35 percent of default rate variation is due to systematic and industry factors. Approximately one third of systematic variation is captured by macroeconomic/financial factors. The remainder is captured by frailty (about 40 percent) and industry (about 25 percent) effects. The default-specific effects are particularly relevant before and during times of financial turbulence. For example, we detect a build-up of systematic risk over the period preceding the 2008 credit crisis.

Keywords: financial crisis; default risk; credit portfolio models; frailty-correlated defaults; state space methods.

JEL classification: C33, G21.

1 Introduction

We decompose default risk of U.S. firms into its different systematic components using a high-dimensional and partly nonlinear non-Gaussian dynamic factor model. Systematic default rate variation, also known as default clustering, constitutes one of the main risks in the banking book of financial institutions. It is well known that corporate default clustering is empirically relevant. For example, aggregate U.S. default rates during the 1991, 2001, and 2008 recession periods are up to five times higher than in intermediate expansion years. It is also well known that default rates depend on the prevailing macroeconomic conditions, see, for example, Pesaran, Schuermann, Treutler, and Weiner (2006), Duffie, Saita, and Wang (2007), Figlewski, Frydman, and Liang (2008), and Koopman, Kräussl, Lucas, and Monteiro (2009).

The common dependence of corporate credit quality on macroeconomic conditions is not the only explanation provided in the literature for default clustering. Recent research indicates that conditioning on readily available macroeconomic and firm-specific information, though important, is not sufficient to fully explain the observed degree of default rate variation. Das, Duffie, Kapadia, and Saita (2007) reject the joint hypothesis of (i) well-specified default intensities in terms of observed macroeconomic and firm-specific information, and (ii) the doubly stochastic independence assumption which underlies many credit risk models that are used in practice. From this finding, two important separate strands of literature have emerged, focusing on frailty-correlated defaults and contagion.

In a frailty model, the additional variation in default intensities is captured by a latent dynamic process, or an unobserved component, see Das et al. (2007), McNeil and Wendin (2007), Koopman, Lucas, and Monteiro (2008), Koopman and Lucas (2008), and Duffie, Eckner, Horel, and Saita (2009). This frailty factor captures default clustering above and beyond what can be explained by macroeconomic variables and firm-specific information.

The unobserved component can capture effects of omitted variables in the model as well as other effects that are difficult to quantify, see Duffie et al. (2009). Contagion models, by contrast, focus on the phenomenon that a defaulting firm weakens other firms with which it has business links, see Giesecke (2004) and Giesecke and Azizpour (2008). Adverse contagion effects may dominate potentially offsetting competitive effects at the intra-industry level, see e.g. Lang and Stulz (1992) and Jorion and Zhang (2007). Lando and Nielsen (2009) screen hundreds of default histories for evidence of direct default contagion. Their results suggest that domino style contagion is a minor concern. Indirect spillover effects, through balance sheet covariates or fire sales, may nevertheless still explain some default dependence at the industry level beyond that induced by shared exposure to macroeconomic/financial (macro) and default-specific (frailty) factors.

It is not known to what extent the three different explanations (macro, frailty, and industry effects) for default clustering interconnect. In particular, it is not yet clear how to measure the relative contribution of the different sources of systematic default risk to observed default clustering. This question is fundamental to our understanding and modeling of default risk. Lando and Nielsen (2009) discuss whether default clustering can be compared with asthma or the flu. In the case of asthma, occurrences are not contagious but depend on exogenous background processes such as air pollution. On the other hand, the flu is directly contagious. Frailty models are, in a sense, more related to models for asthma, while contagion models based on self-exciting processes are similar to models for flu. Whether one effect dominates the other empirically is therefore highly relevant to the appropriate modeling framework for portfolio credit risk.

We decompose the systematic variation in corporate defaults into different constituents within a high-dimensional and partly nonlinear non-Gaussian dynamic factor model. Within this modeling framework, we let default rate volatility at the rating and industry level be attributed to the macro, frailty, and industry effects, simultaneously. The estimation of

these dynamic factors and the other parameters in the model is carried out by Monte Carlo maximum likelihood methods. The implementation details of the estimation methods are presented. The attractive feature of our framework is threefold. First, it allows us to combine typical Gaussian time series (such as macroeconomic variables, business cycle indicators, financial market conditions, and interest rates) with discrete time series such as default counts. Second, and in contrast to earlier models, we can include a substantive number of macroeconomic/financial variables to account for the macroeconomic conditions. Third, our new framework allows for an integrated view on the interaction between macro, frailty, and industry factors by treating them simultaneously rather than in a typical two-step estimation approach. This proves to be very convenient if the empirical model is also used for computing adequate economic capital buffers and in a stress testing exercises.

Our estimation results indicate that defaults are more related to asthma than to flu: the common factors to all firms (macro and frailty) account for approximately 75% of the default clustering. It leaves industry (and thus possibly contagion) effects as a substantial secondary source of credit portfolio risk. To quantify these contributions to systematic default risk, we introduce a pseudo- R^2 measure of fit based on reductions in Kullback-Leibler (KL) divergence. The KL divergence is a standard statistical measure of ‘distance’ between distributions and reduces to the usual R^2 in a linear regression model. Its use is appropriate in a context where there are both discrete (default counts) and continuous (macro variables) data. We find that on average across industries and time, 66% of total default risk is idiosyncratic and therefore diversifiable. The remainder 34% is systematic. For subinvestment grade firms, one third of systematic default risk can be attributed to common variation with the business cycle and with financial markets data. For investment grade firms, this percentage is as high as 60%. The remaining share of systematic credit risk is driven by a frailty factor and industry-specific factors (in approximately equal proportions). The frailty component cannot be diversified in the cross-section, whereas the industry effects

can only be diversified to some extent.

Our reported risk shares vary considerably over industry sectors, rating groups and time. For example, we find that the frailty component tends to explain a higher share of default rate volatility before and during times of crisis. In particular, we find systematic credit risk building up in the years 2002-2008, leading up to the financial crisis, when default activity was much lower than suggested by macro-financial data. The framework may thus also provide a tool to detect systemic risk build-up in the economy. Tools to assess the evolution and composition of latent financial risks are urgently needed at macro-prudential policy institutions, such as the Financial Services Oversight Council (FSOC) in the United States, and the European Systemic Risk Board (ESRB) in the European Union.

In this study we also seek to address the question which missing sources of default rate volatility the frailty factor may be capturing. Interestingly, we find a positive correlation between changes in our estimated frailty factor and proxies for tightening lending standards. This, along with other pieces of evidence, suggests that the frailty component may be able to capture changes in credit supply conditions and the ease of credit access. Credit supply and credit risk are clearly connected: it is hard to default if available credit is plentiful. Conversely, even solvent firms can get into financial distress if credit is heavily rationed. Changes in the ease of credit access are typically hard to quantify empirically as they may depend on many developments which are also hard to measure, such as changing activity in the securitization market or changes in banks' business models. Still, when combined all these factors may have a systematic and economically significant impact on the loss experience of diversified credit portfolios.

In relation to the 2008 crisis period, the following findings from our empirical study are relevant. First, significant frailty effects imply that default and business cycle do not coincide. They have diverged significantly and persistently in the past, and most recently during the run-up to the 2008 credit crisis. Such a decoupling may indicate a credit bubble, in

particular if in addition credit quantity growth is unusually high and bank lending standards are low. Second, stressing the usual macro-financial covariates may not be sufficient to assess financial stability conditions in a stress test. Systematic default rate increases also depend on additional latent systematic risk factors. These may need to be stressed just as the observed systematic risk factors are stressed. An admittedly incomplete understanding of latent credit risk sources does not change the fact that they matter empirically. Finally, while there is a long tradition in central banks of analyzing credit *quantities* over time and comparing them to fundamentals (such as tracking the private-credit to GDP ratio), the tracking of credit *risk* conditions and its composition has received less to no attention at all. The new econometric methodology developed in this study is a versatile framework for making the latter operational.

The remainder of this paper is organized as follows. Section 2 introduces our general methodological framework. Section 3 presents our core empirical results, in particular a decomposition of total systematic default risk into its latent constituents. We comment on implications for portfolio credit risk in Section 4. Section 5 concludes.

2 A joint model for default, macro, and industry risk

The key challenge in decomposing systematic credit risk is to define a factor model structure that can simultaneously handle normally distributed (macro variables) and non-normally distributed (default counts) data, as well as linear and non-linear factor dependence. The factor model we introduce for this purpose is a Mixed Measurement Dynamic Factor Model, or in short, MiMe DFM. In the development of our new model, we focus on the decomposition of systematic default risk. However, the model may also find relevant applications in other areas of finance. The model is applicable to any setting where different distributions have to be mixed in a factor structure.

In our analysis we consider the vector of observations given by

$$y_t = (y_{1t}, \dots, y_{Jt}, y_{J+1,t}, \dots, y_{J+N,t})', \quad (1)$$

with time index $t = 1, \dots, T$. The first J elements of y_t are default counts. We count defaults for different rating groups and industries. As a consequence, the first J elements of y_t contain discrete, non-negative values. The remaining N elements of y_t contain macro and financial variables which are typically taken as Gaussian variables. We assume that the default counts and the macro and financial time series data are subject to a set of dynamic factors. Some of these factors may be common to all variables in y_t . Other factors may only affect a subset of the elements in y_t .

2.1 The mixed measurement dynamic factor model

We distinguish macro, frailty, and industry factors; these common factors are denoted as f_t^m , f_t^d , and f_t^i , respectively. The factors f_t^m capture shared business cycle dynamics in macroeconomic variables and default counts. Therefore, factors f_t^m are common to all variables in y_t . Common frailty factors f_t^d are default-specific, i.e., common to default counts (y_{1t}, \dots, y_{Jt}) only and independent of observed macroeconomic and financial data by construction. By not allowing the frailty factors to impact the macro series y_{jt} for $j = J + 1, \dots, J + N$, we effectively restrict f_t^d to capture default clustering above and beyond that is implied by macroeconomic and financial factors f_t^m . The third set of factors f_t^i affects firms in the same industry. Such factors may be particularly relevant in specific industries such as the financial industry. In addition, our estimated industry frailty factors may partly capture contagion driven default clustering in specific industries.

We gather all factors into the vector $f_t = (f_t^{m'}, f_t^{d'}, f_t^{i'})'$ which we model as a dynamic latent (unobserved) vector variable. We adopt an autoregressive dynamic process for the

latent factors,

$$f_t = \Phi f_{t-1} + \eta_t, \quad t = 1, 2, \dots, T, \quad (2)$$

where the coefficient matrix Φ is assumed a diagonal matrix and with the $m \times 1$ disturbance vector $\eta_t \sim \text{NID}(0, \Sigma_\eta)$ which is serially uncorrelated. The standard stationarity conditions apply to f_t . To complete the specification of the factor process in (2), we specify the initial factor by $f_1 \sim \text{N}(0, \Sigma_f)$ where Σ_f is the unconditional variance matrix of f_t and the solution of the matrix equation $\Sigma_f = \Phi \Sigma_f \Phi' + \Sigma_\eta$.

More elaborate dynamic processes for f_t can also be considered. The autoregressive structure in (2) allows the components of f_t to be sticky. The macroeconomic factors f_t^m evolve slowly over time and capture business cycle effects in both macro and default data. The credit climate and industry default conditions are represented by persistent processes for f_t^d and f_t^i which typically capture the clustering of defaults during high-default years.

Conditional on f_t , the first J elements of y_t are modeled as binomial densities with parameters k_{jt} and π_{jt} , for $j = 1, \dots, J$, where k_{jt} is the number of firms and π_{jt} is the probability of default in the industry and rating group j at time t . Exposures k_{jt} are counted at the beginning of each time period t , and are held fixed during this period. For more details on the conditionally binomial model, see e.g. McNeil, Frey, and Embrechts (2005, Chapter 9). Frey and McNeil (2002) show that almost all available industry credit risk models, such as Creditmetrics, Moody's KMV, and CreditRisk+ can be presented as conditional binomial models. The last N elements of y_t are, conditional on f_t , distributed as independently normal variables with mean μ_{jt} and fixed variance σ_j^2 for $j = J+1, \dots, J+N$.

The density specifications for the default counts and the macro variables can now be given by

$$\begin{aligned} y_{jt}|f_t &\sim \text{Bin}(\pi_{jt}, k_{jt}), & \text{for } j = 1, \dots, J, \\ y_{jt}|f_t &\sim \text{N}(\mu_{jt}, \sigma_j^2), & \text{for } j = J+1, \dots, J+N. \end{aligned} \quad (3)$$

where Bin refers to the Binomial density and N to the normal density. The support of

probability π_{jt} is restricted between zero and one. We enforce this by considering the logit transformation

$$\pi_{jt} = \frac{\exp \pi_{jt}^*}{1 + \exp \pi_{jt}^*}, \quad j = 1, \dots, J.$$

The auxiliary variable π_{jt}^* and the mean of the normal density μ_{jt} are linear functions of the factor f_t and given by

$$\pi_{jt}^* = \lambda_j + \beta_j' f_t^m + \gamma_j' f_t^d + \delta_j' f_t^i, \quad \text{for } j = 1, \dots, J, \quad (4)$$

$$\mu_{jt} = \lambda_j + \beta_j' f_t^m, \quad \text{for } j = J + 1, \dots, J + N, \quad (5)$$

where λ_j is a constant and β_j , γ_j and δ_j are column vectors of loading coefficients with appropriate dimensions. The number of firms k_{jt} is known. The variance σ_j^2 is assumed fixed.

In this particular specification of the MiMe DFM, we can measure the relative contributions of macro, frailty, and industry risk to general default risk. The factors in f_t^m capture general developments such as business cycle activity, lending conditions and developments in financial markets. The auxiliary variable for default probability in (4) partly depends on macro factors, but also depends on frailty risk f_t^d and industry f_t^i factors. The specifications in (4) and (5) are key to our empirical analysis where we focus on studying whether macro dynamics explain all systematic default rate variation, or whether and to what extent frailty and industry factors are also important.

The estimation of the constants λ_j and the factor loadings β_j , $\gamma_{j'}$ and $\delta_{j'}$, for $j = 1, \dots, J + N$ and $j' = 1, \dots, J$, together with the diagonal elements of Φ in (2) is carried out by the method of Monte Carlo maximum likelihood, see Section 2.2. For the identification of the factor loadings, we require standardized factors f_t . We therefore restrict the variance matrix $\Sigma_f = I$ and hence the variance matrix of the disturbance vector η_t in (2) becomes $\Sigma_\eta = I - \Phi\Phi'$. The estimation of the variances σ_j^2 can be circumvented since we assume that our macroeconomic variables are standardized. This is common practice in the macroeconomic

forecasting literature, see e.g. Stock and Watson (2002). We then have $\text{Var}(y_{jt}) = \beta_j' \Sigma_f \beta_j + \sigma_j^2 = 1$ and hence $\sigma_j^2 = 1 - \beta_j' \beta_j$ since we have assumed that $\Sigma_f = I$.

2.2 Parameter estimation via importance sampling

An analytical expression for the the maximum likelihood (ML) estimate of parameter vector ψ for the MiMe DFM is not available. A feasible approach to the ML estimation of ψ is the maximization of the likelihood function that is evaluated via Monte Carlo methods such as importance sampling. A short description of this approach is given below. A full treatment is presented by Durbin and Koopman (2001, Part II).

The observation density function of $y = (y'_1, \dots, y'_T)'$ can be expressed by the joint density of y and $f = (f'_1, \dots, f'_T)'$ where f is integrated out, that is

$$p(y; \psi) = \int p(y, f; \psi) df = \int p(y|f; \psi) p(f; \psi) df, \quad (6)$$

where $p(y|f; \psi)$ is the density of y conditional on f and $p(f; \psi)$ is the density of f . A Monte Carlo estimator of $p(y; \psi)$ can be obtained by

$$\hat{p}(y; \psi) = M^{-1} \sum_{k=1}^M p(y|f^{(k)}; \psi), \quad f^{(k)} \sim p(f; \psi),$$

for some large integer M . The estimator $\hat{p}(y; \psi)$ is however numerically inefficient since most draws $f^{(k)}$ will not contribute substantially to $p(y|f; \psi)$ for any ψ and $k = 1, \dots, K$. Importance sampling improves the Monte Carlo estimation of $p(y; \psi)$ by sampling f from the Gaussian importance density $g(f|y; \psi)$. We can express the observation density function $p(y; \psi)$ by

$$p(y; \psi) = \int \frac{p(y, f; \psi)}{g(f|y; \psi)} g(f|y; \psi) df = g(y; \psi) \int \frac{p(y|f; \psi)}{g(y|f; \psi)} g(f|y; \psi) df. \quad (7)$$

Since f is from a Gaussian density, we have $g(f; \psi) = p(f; \psi)$ and $g(y; \psi) = g(y, f; \psi) / g(f|y; \psi)$.

In case $g(f|y; \psi)$ is close to $p(f|y; \psi)$ and in case simulation from $g(f|y; \psi)$ is feasible, the

Monte Carlo estimator of the likelihood function is given by

$$\tilde{p}(y; \psi) = g(y; \psi) M^{-1} \sum_{k=1}^M \frac{p(y|f^{(k)}; \psi)}{g(y|f^{(k)}; \psi)}, \quad f^{(k)} \sim g(f|y; \psi), \quad (8)$$

is numerically much more efficient, see Kloek and van Dijk (1978), Geweke (1989) and Durbin and Koopman (2001).

The importance density $g(f|y; \psi)$ is based on an approximating, linear Gaussian state space model based on an observation equation for each y_{jt} in (1) and given by

$$y_{jt} = c_{jt} + \theta_{jt} + \varepsilon_{jt}, \quad \varepsilon_{jt} \sim N(0, h_{jt}), \quad (9)$$

where c_{jt} is a known mean, θ_{jt} is the unobserved signal and h_{jt} is a known variance, for $j = 1, \dots, J + N$. For the normal variables y_{jt} , the signal θ_{jt} is equal to μ_{jt} of (5) and the variables $c_{jt} = 0$ and $h_{jt} = \sigma_j^2$ are known with $j = J + 1, \dots, J + N$. For the default counts y_{jt} in the approximating model, we let the signal θ_{jt} be equal to π_{jt}^* of (4), with $j = 1, \dots, J$. The variables c_{jt} and h_{jt} for the default counts are determined such that the modes of $p(f|y; \psi)$ and $g(f|y; \psi)$ are equal, see Shephard and Pitt (1997) and Durbin and Koopman (1997) for the details. The values for c_{jt} and h_{jt} are found iteratively and by means of the Kalman filter and an associated smoothing method.

To simulate values from the resulting importance density $g(f|y; \psi)$ based on the approximating model (9), the simulation smoothing method of Durbin and Koopman (2002) can be used. For a set of M draws $f^{(1)}, \dots, f^{(M)}$ from $g(f|y; \psi)$, the evaluation of the likelihood function (8) via importance sampling relies on the computation of $p(y|f; \psi)$, $g(y|f; \psi)$, with $f = f^{(k)}$, and $g(y; \psi)$ for $k = 1, \dots, M$. Density $p(y|f; \psi)$ is based on the model specifications in (3). Density $g(y|f; \psi)$ is based on the approximating, linear Gaussian model (9). Density $g(y; \psi)$ is effectively the likelihood function of the approximating model (9) and can be computed via the Kalman filter, see Durbin and Koopman (2001). Testing the assumptions underlying the application of importance sampling can be carried out using the procedures proposed by e.g. Koopman, Shephard, and Creal (2009).

2.3 Estimation of latent factors

Inference on the latent factors can also be based on importance sampling. In particular, it can be shown that

$$E(f|y; \psi) = \int f \cdot p(f|y; \psi) df = \frac{\int f \cdot w(y, f; \psi) g(f|y; \psi) df}{\int w(y, f; \psi) g(f|y; \psi) df},$$

where $w(y, f; \psi) = p(y|f; \psi)/g(y|f; \psi)$. The estimation of $E(f|y; \psi)$ via importance sampling can be achieved by

$$\tilde{f} = \frac{\sum_{k=1}^M w_k \cdot f^{(k)}}{\sum_{k=1}^M w_k},$$

with $w_k = p(y|f^{(k)}; \psi)/g(y|f^{(k)}; \psi)$ and where $f^{(k)} \sim g(f|y; \psi)$ is obtained by simulation smoothing. The standard error of \tilde{f}_i , the i th element of \tilde{f} , is denoted by s_i and is computed by

$$s_i^2 = \left(\frac{\sum_{k=1}^M w_k \cdot (f_i^{(k)})^2}{\sum_{k=1}^M w_k} \right) - \tilde{f}_i^2,$$

where $f_i^{(k)}$ is the i th element of $f^{(k)}$.

2.4 Decomposition of the default count variation

Given the estimated parameters and risk factors, we may wish to assess which share of variation in default data is captured by the different risk factors. For this purpose we adopt a pseudo- R^2 measure that is discussed in Cameron and Windmeijer (1997). It is based on the Kullback-Leibler divergence measure and is defined as

$$KL(\theta_1, \theta_2) = 2 \int [\log p_{\theta_1}(y) - \log p_{\theta_2}(y)] p_{\theta_1}(y) dy \quad (10)$$

where $p_{\theta_i}(y)$ is the density of the model with signal vector $\theta = \theta_i$ for $i = 1, 2$. The signal vector θ refers to all π_{jt}^* in (4) and $\mu_{j't}$ in (5) for $j = 1, \dots, J$, $j' = J + 1, \dots, J + N$ and $t = 1, \dots, n$. The vectors θ_1 and θ_2 refer to signals which are composed of different selections of factors f_t^m , f_t^d and f_t^i in (4) and (5). The $KL(\theta_1, \theta_2)$ divergence in (10) measures the

distance between the log-densities $\log p_{\theta_1}$ and $\log p_{\theta_2}$ with respect to the density of the model with signal θ_1 . For example, if the densities in (10) are normal, $KL(\theta_1, \theta_2)$ measures the increase in the sum of squared residuals of model with θ_2 in relation to model with θ_1 . In our set of models, $p_{\theta}(y)$ refers to both normal and binomial distributions.

The pseudo- R^2 is defined as the proportional reduction in variation of default rates due to the inclusion of additional factors, that is

$$R^2(\theta) = 1 - \frac{KL(\theta_{max}, \theta)}{KL(\theta_{max}, \theta_{na})}, \quad (11)$$

where $KL()$ is defined in (10). The signal θ_{na} does not depend on any factor and consists of the constants λ_j only, for $j = 1, \dots, J + N$. The model with signal θ_{max} provides the maximum possible fit as it contains a separate dummy variable for each observation. The model contains as many parameters as observations. While this unrestricted model is not useful for practical purposes, it does provide a benchmark for the maximum possible fit. Figure 1 illustrates the KL measures from which we can compute the pseudo- R^2 measures. We distinguish several alternative model specifications indicated by their signals θ_{na} , θ_m , θ_{md} , and θ_{mdi} which contain an increasing collection of latent factors. The models with θ_m , θ_{md} , and θ_{mdi} cumulate the macro f_t^m , frailty f_t^d , and industry f_t^i factors, respectively.


The value of $R^2(\theta)$ is by construction between zero and one. The relative contribution from each of our systematic credit risk factors is measured as the increase in the pseudo- R^2 value when moving from θ_m via θ_{md} to θ_{mdi} . The remainder increase from θ_{mdi} to θ_{max} can be qualified as idiosyncratic risk.

3 Empirical findings for U.S. default and macro data

We study the quarterly default and exposure counts obtained from the Moody's corporate default research database for the period 1971Q1 to 2009Q1. We distinguish seven industry groups (financials and insurance; transportation; media, hotels, and leisure; utilities and

Figure 1: Models and reductions in the Kullback-Leibler divergence

The graph shows how reductions in the estimated KL divergence are used to decompose the total variation in non-Gaussian default counts into risk shares corresponding to models with increasing sets of latent factors (M^{na} , M^m , M^{md} , $M^{m di}$ and M^{max}).



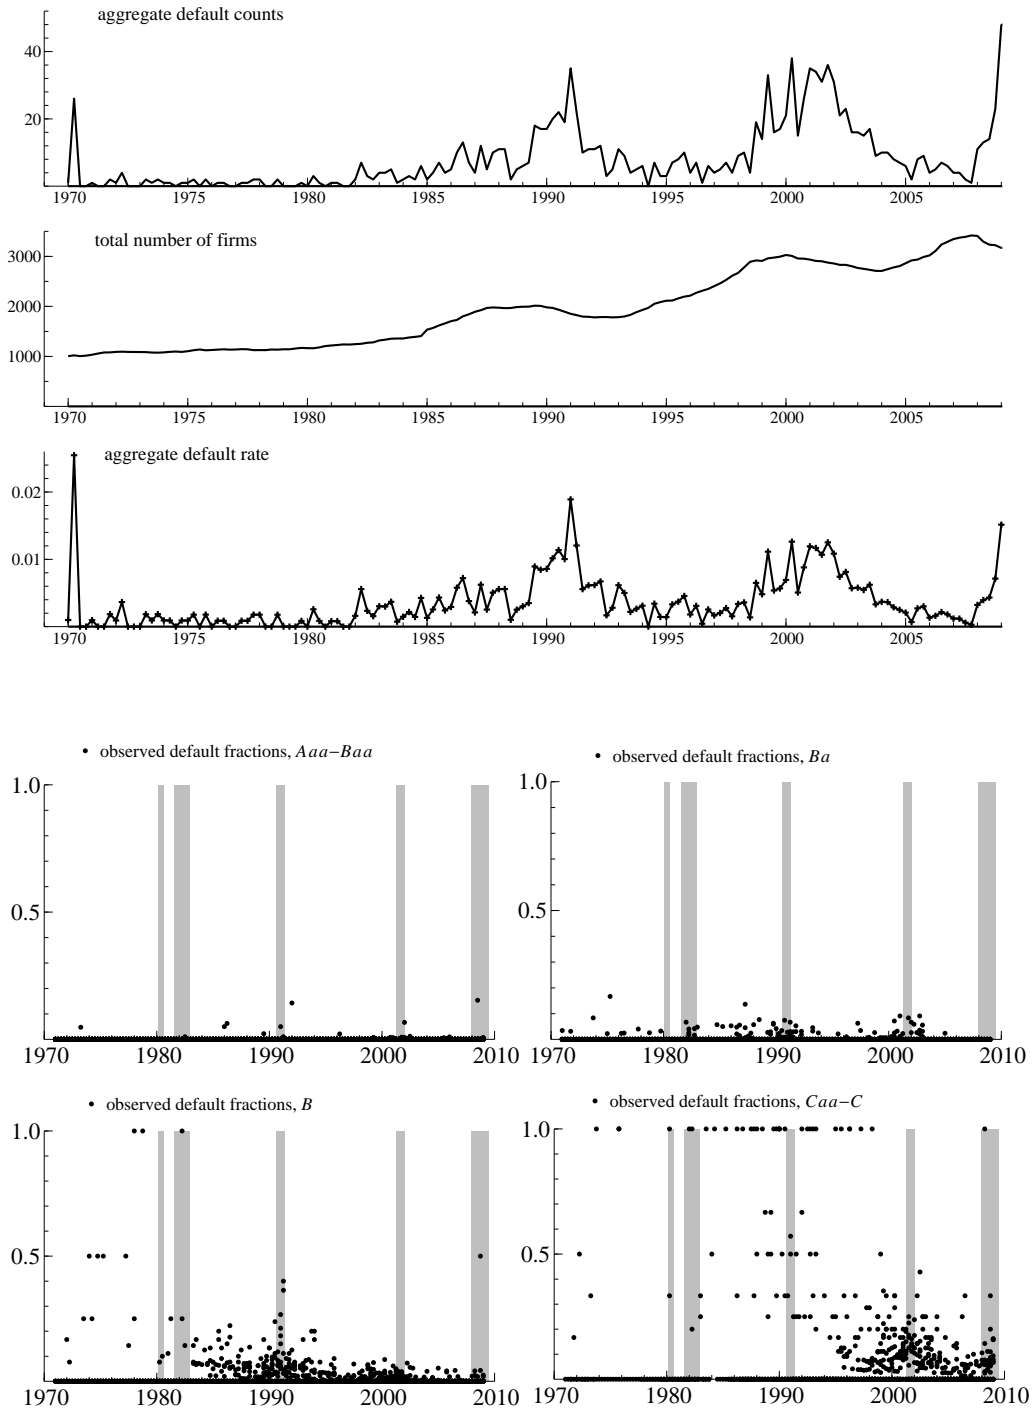
energy; industrials; technology; and retail and consumer products) and four rating groups (investment grade $Aaa - Baa$, and the speculative grade groups $Ba, B, Caa - C$). We have pooled the investment grade firms because defaults are rare for this segment. It is assumed that current issuer ratings summarize the available information about a firm's financial strength. This may be true only to a first approximation. However, rating agencies take into account a vast number of accounting and management information, and provide an assessment of firm-specific information which is comparable across industry sectors. While we focus on issuer ratings as available to us from Moody's, ratings could alternatively be constructed from mapping EDF and CDS data into rating bins where available. In these cases, however, limited sample sizes are typically an issue. Finally, other factors may help in explaining firm specific default risks, such as distance-to-default measures, trailing stock returns and accounting information (leverage). However, improving risk prediction for single-name credit risks is not the aim of this study. Ratings are taken into account mainly because macro, frailty, and industry effects (in terms of factor loadings) may be different for financially healthy and financially weaker firms. Ratings are likely to be sufficiently accurate for that purpose.

Figure 2 presents aggregate default fractions and disaggregated default data. We observe a considerable time variation in aggregate default fractions. The disaggregated data reveals that defaults cluster around recession periods for both investment grade and speculative grade rated firms.

Macroeconomic and financial data are obtained from the St. Louis Fed online database FRED, see Table 1 for a listing of macroeconomic and financial data. The panel data are available on a monthly basis. We consider macro-financial covariates that are typically also stressed in a macro stress test, as indicated in for example CEBS (2010) and Tarullo (2010). These usually involve business cycle measurements (production and income), labor market conditions (unemployment rate), short and long term interest rates (term structure) and

Figure 2: Clustering in default data

The first three panels present time series of (i) the total number of defaults in the Moody's database $\sum_j y_{jt}$, (ii) the total number of exposures $\sum_j k_{jt}$, and (iii) the aggregate default rate for all Moody's rated U.S. firms, $\sum_j y_{jt} / \sum_j k_{jt}$. The bottom four graphs present the observed default fractions y_{jt}/k_{jt} over time. We distinguish four broad rating groups: *Aaa – Baa*, *Ba*, *B*, and *Caa – C*. Each such panel plots disaggregated time series of industry-specific default fractions (which are mostly zero). Shaded areas correspond to NBER U.S. recession periods.



credit spreads, as well as stock market returns and volatilities. The data enters the analysis in the form of annual growth rates, see Figure 3 for time series plots.

The panel information criteria based on Bai and Ng (2002) are standard measures to determine the number of factors in large factor models. These criteria have been adapted by Alessi, Barigozzi, and Capasso (2008) to improve robustness for panels of smaller cross sectional dimensions. According to these information criteria, two to four macro factors are appropriate to summarize the information in the macro panel. We therefore include four macro factors in our final specification. Allowing for one frailty factor is standard in the literature, see McNeil and Wendin (2007), Duffie et al. (2009) and Azizpour et al. (2010), and appears sufficient to capture deviations of credit from macro conditions when modeling defaults. Finally, we allow for six industry-specific factors driving defaults for firms in certain broad industry groups.

Table 1: Macroeconomic Time Series Data

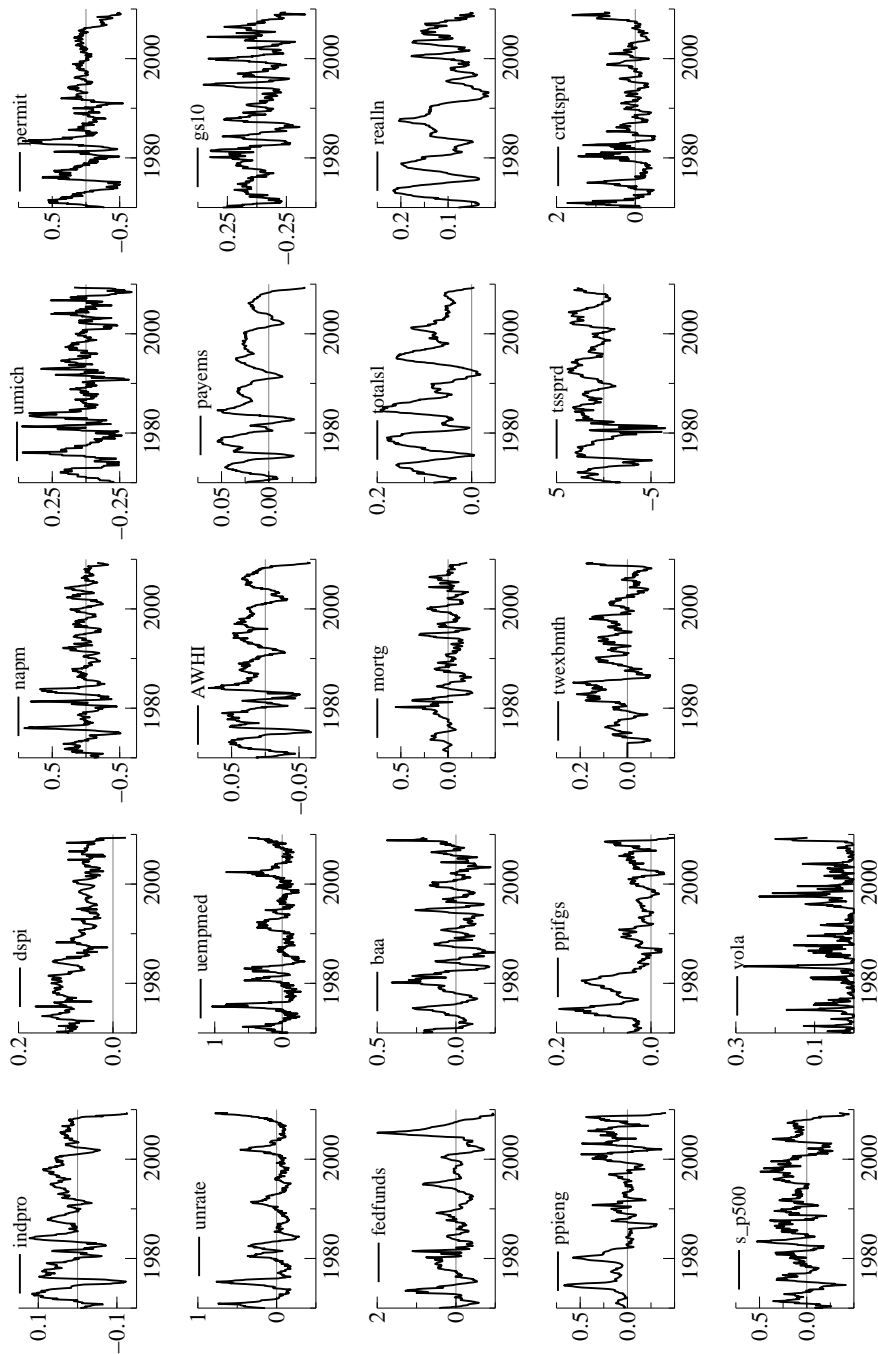
The table gives a full listing of included macroeconomic time series data x_t and binary indicators b_t . All time series are obtained from the St. Louis Fed online database, <http://research.stlouisfed.org/fred2/>.

Category	Summary of time series in category	Shortname	Total no
(a) Macro indicators, and business cycle conditions	Industrial production index	indpro	5
	Disposable personal income	dspi	
	ISM Manufacturing index	napm	
	Uni Michigan consumer sentiment	umich	
	New housing permits	permit	
(b) Labour market conditions	Civilian unemployment rate	unrate	4
	Median duration of unemployment	uempmed	
	Average weekly hours index	AWHI	
	Total non-farm payrolls	payems	
(c) Monetary policy and financing conditions	Government bond term structure spread	gs10	6
	Federal funds rate	fedfunds	
	Moody's seasoned Baa corporate bond yield	baa	
	Mortgage rates, 30 year	mortg	
	10 year treasury rate, constant maturity	tssprd	
	Credit spread corporates over treasuries	creditsprd	
(d) Bank lending	Total Consumer Credit Outstanding	totalsl	2
	Total Real Estate Loans, all banks	realln	
(e) Cost of resources	PPI Fuels and related Energy	ppieng	3
	PPI Finished Goods	ppifgs	
	Trade-weighted U.S. dollar exchange rate	twexbmth	
(f) Stock market returns	S&P 500 yearly returns	s_p500	2
	S&P 500 return volatility	vola	

22

Figure 3: Macroeconomic and financial time series data

We present time series of yearly growth rates in macroeconomic and financial data. For a listing of the data we refer to Table 1.



3.1 Parameter and risk factor estimates

Parameter estimates associated with the default counts are presented in Table 2. Estimated coefficients refer to a model specification with macroeconomic, frailty, and industry-specific factors. Parameter estimates in the first column combine to fixed effects for each cross-section j , according to $\lambda_j = \lambda_0 + \lambda_{1,r_j} + \lambda_{2,s_j}$, where the common intercept λ_0 is adjusted by specific coefficients indicating industry sector (s_j) and rating group (r_j), respectively, for $j = 1, \dots, J$ with J as the total number of unique groups. The second column reports the factor loadings β associated with four common macro factors f_t^m . Loading coefficients differ across rating groups. The loadings tend to be larger for investment grade firms; in particular, their loadings associated with macro factors 1, 3, and 4 are relatively large. This finding confirms that financially healthy firms tend to be more sensitive to business cycle risk, see e.g. Basel Committee on Banking Supervision (2004).

Factor loadings γ and δ are given in the last two columns of Table 2. The loadings in γ are associated with a single common frailty factor f_t^d while the loadings in δ are for the six orthogonal industry factors f_t^i . The frailty risk factor f_t^d is, by construction, common to all firms, but unrelated to the macroeconomic data. Frailty risk is relatively large for all firms, but particularly pronounced for speculative grade firms. Industry sector loadings are highest for the financial, transportation, and energy and utilities sector.

The top panel of Figure 4 presents the estimated risk factors f_t^m as defined in (4) and (5). We plot the estimated conditional mean of the factors, along with approximate standard error bands at a 95% confidence level. For estimation details, we refer to the Appendix. The factors are ordered row-wise from top-left to bottom-right according to their share of explained variation for the macro and financial data listed in Table 1.

The bottom panel of Figure 4 presents the shares of variation in each macroeconomic time series that can be attributed to the common macroeconomic factors. The first two

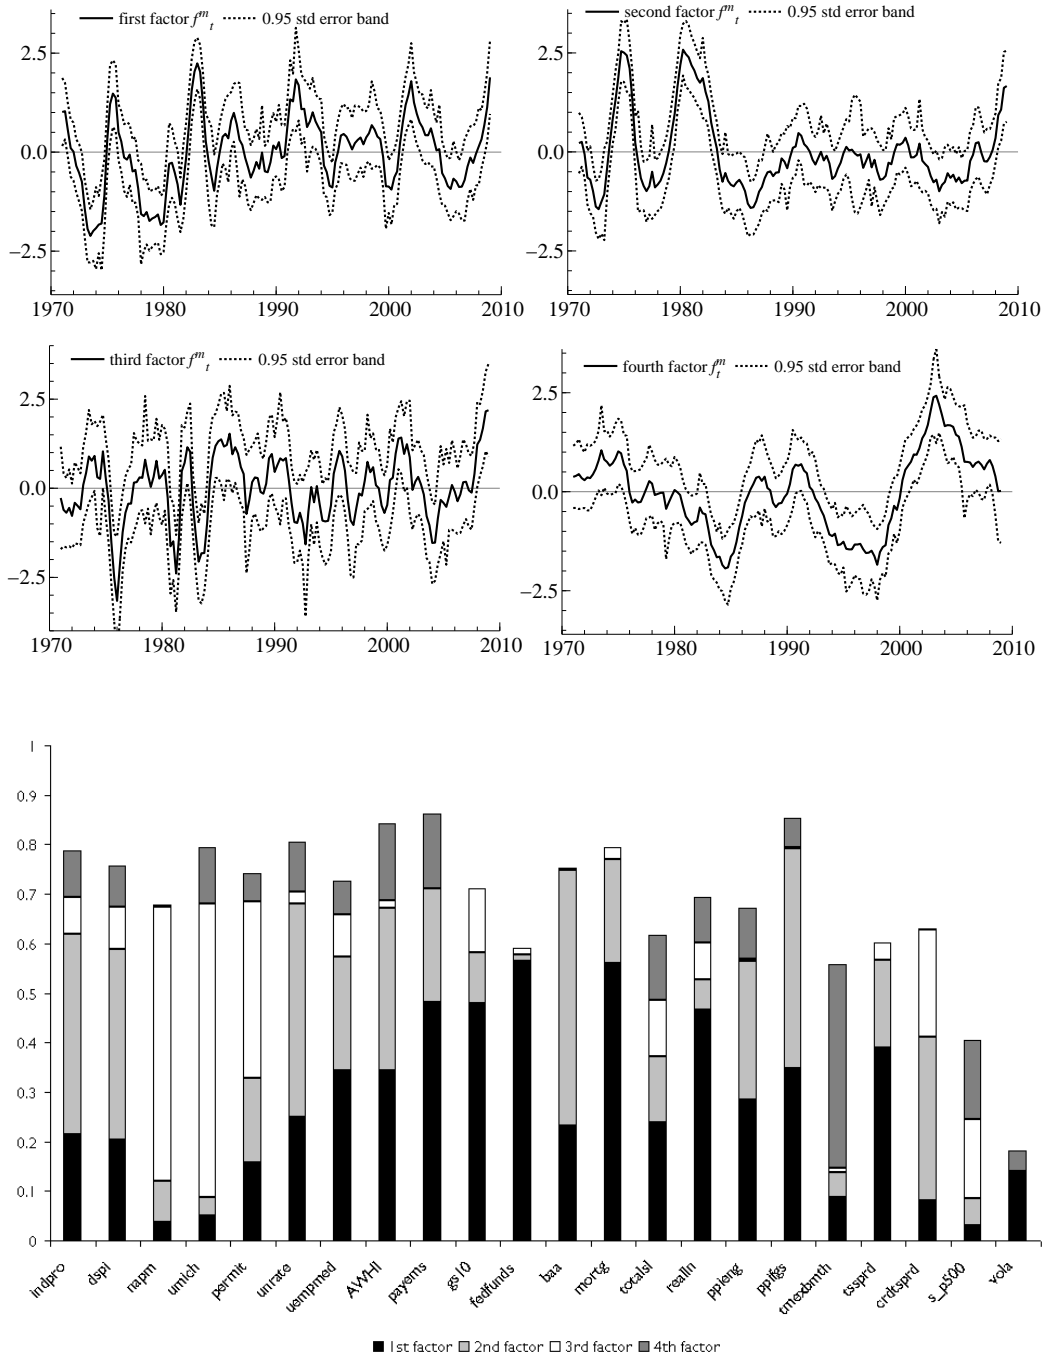
Table 2: Parameter estimates (for the binomial series)

We report parameter estimates associated with the binomial part of our model. The coefficients in the first column combine to fixed effects according to $\lambda_j = \lambda_0 + \lambda_{1,r_j} + \lambda_{2,s_j}$, the common intercept λ_0 is adjusted to take into account a fixed effect for the rating group and industry sector. The middle column reports autoregressive coefficients ϕ_k (the k th diagonal element of Φ in (2) for $k = 1, 2, 3, 4$) and loading coefficients in β_j on four common macro factors f_t^m . The last column reports the loading coefficients γ_j on the frailty factor f_t^d and loadings δ_j on industry-specific risk factors f_t^i . The estimation sample is from 1971Q1 to 2009Q1.

Intercepts λ_j			Loadings f_t^m			Loadings f_t^d		
par	val	t-val	par	val	t-val	par	val	t-val
λ_0	-2.52	7.61	ϕ_1	0.89	2.61	γ_{IG}	0.16	0.75
			$\beta_{1,IG}$	0.57	1.05	γ_{Ba}	0.52	2.15
λ_{fin}	-0.21	0.82	$\beta_{1,Ba}$	0.30	0.55	γ_B	0.72	4.80
λ_{tra}	-0.02	0.07	$\beta_{1,B}$	0.26	0.72	γ_C	0.43	5.92
λ_{lei}	-0.17	0.75	$\beta_{1,C}$	0.17	0.57			
λ_{utl}	-0.76	2.02				Loadings f_t^i		
λ_{tec}	-0.09	0.64	ϕ_2	0.91	2.21	δ_{fin}	0.73	2.96
λ_{ret}	-0.32	1.68	$\beta_{2,IG}$	-0.01	-0.06	δ_{tra}	0.61	1.85
			$\beta_{2,Ba}$	0.05	0.10	δ_{lei}	0.41	2.62
λ_{IG}	-6.95	15.30	$\beta_{2,B}$	0.12	0.39	δ_{utl}	0.98	3.87
λ_{BB}	-3.92	14.11	$\beta_{2,C}$	0.24	0.97	δ_{tec}	0.41	2.56
λ_B	-2.22	11.21				δ_{ret}	0.40	2.71
			ϕ_3	0.77	1.64			
			$\beta_{3,IG}$	0.70	1.86			
			$\beta_{3,Ba}$	0.34	0.86			
			$\beta_{3,B}$	0.24	1.08			
			$\beta_{3,C}$	0.18	1.45			
			ϕ_4	0.94	2.99			
			$\beta_{4,IG}$	0.53	0.80			
			$\beta_{4,Ba}$	0.41	0.65			
			$\beta_{4,B}$	-0.08	-0.43			
			$\beta_{4,C}$	0.14	0.37			

Figure 4: Macroeconomic risk factor estimates

The first four panels present the estimated risk factors f_t^m as defined in (4) and (5). We present the estimated conditional mean of the factors, along with approximate standard error bands at a 95% confidence level. We refer to the Appendix for the estimation and signal extraction methodology. The bottom panel indicates which share of the variation in each time series listed in Table 1 can be attributed to each factor f^m . Factors f^m are common to the (continuous) macro and financial as well as the (discrete) default count data.



macroeconomic factors load mostly on labor market, production, and interest rate data. The last two factors displayed in the top panel of Figure 4 load mostly on survey sentiment data and changes in price level indicators. The macroeconomic factors capture 27.2%, 21.3%, 11.7%, and 8.3% of the total variation in the macro data panel, respectively (68.6% in total). All four common factors f_t^m tend to load more on default probabilities of firms rated investment grade rather than speculative grade, see Table 2.

Figure 5 presents conditional mean estimates of the frailty and industry-specific factors. The frailty factor is high *before* and *during* the recession years 1991 and 2001. As a result, the frailty factor implies additional default clustering in these times of stress. On the other hand, the large negative values before the 2007-2009 credit crisis imply defaults that are systematically ‘too low’ compared to what is implied by macroeconomic and financial data. The frailty factor reverts to its mean level during the 2007-09 credit crisis. Apparently, the extreme realizations in macroeconomic and financial variables during 2008-09 are sufficient to account for the levels of observed defaults.

Both Das et al. (2007) and Duffie, Eckner, Horel, and Saita (2009) ask what effects are captured by the frailty factor. Our estimate in Figure 5 suggests that the frailty factor may partly capture the outward shift in credit supply due to a high level of asset securitization activity and a lowering of lending standards during 2005-2007. Conversely, in 2001 and 2002, the frailty factor may capture adverse shocks to credit supply due to the disappearance of trust in accounting information in response to the Enron and Worldcom scandals. These credit supply effects are likely to be important for defaults, but also difficult to measure because only the intersection of credit supply and demand is observed.

We now present two additional pieces of evidence for the claim that our estimated frailty effects in part capture changes in the ease of credit access for credit constrained firms. First, Figure 6 presents the estimated composite default signals θ_{jt} for investment grade industrial firms (Aaa-Baa) and respective speculative grade firms (Ba-C). For investment grade firms,

Figure 5: Frailty risk factor and industry-group dynamics

The first panel presents the estimated conditional mean of the frailty risk factor. This risk factor is common to all default counts. The final six panels present six industry-specific risk factors along with asymptotic standard error bands at a 95% confidence level. High risk factor values imply higher expected default rates.

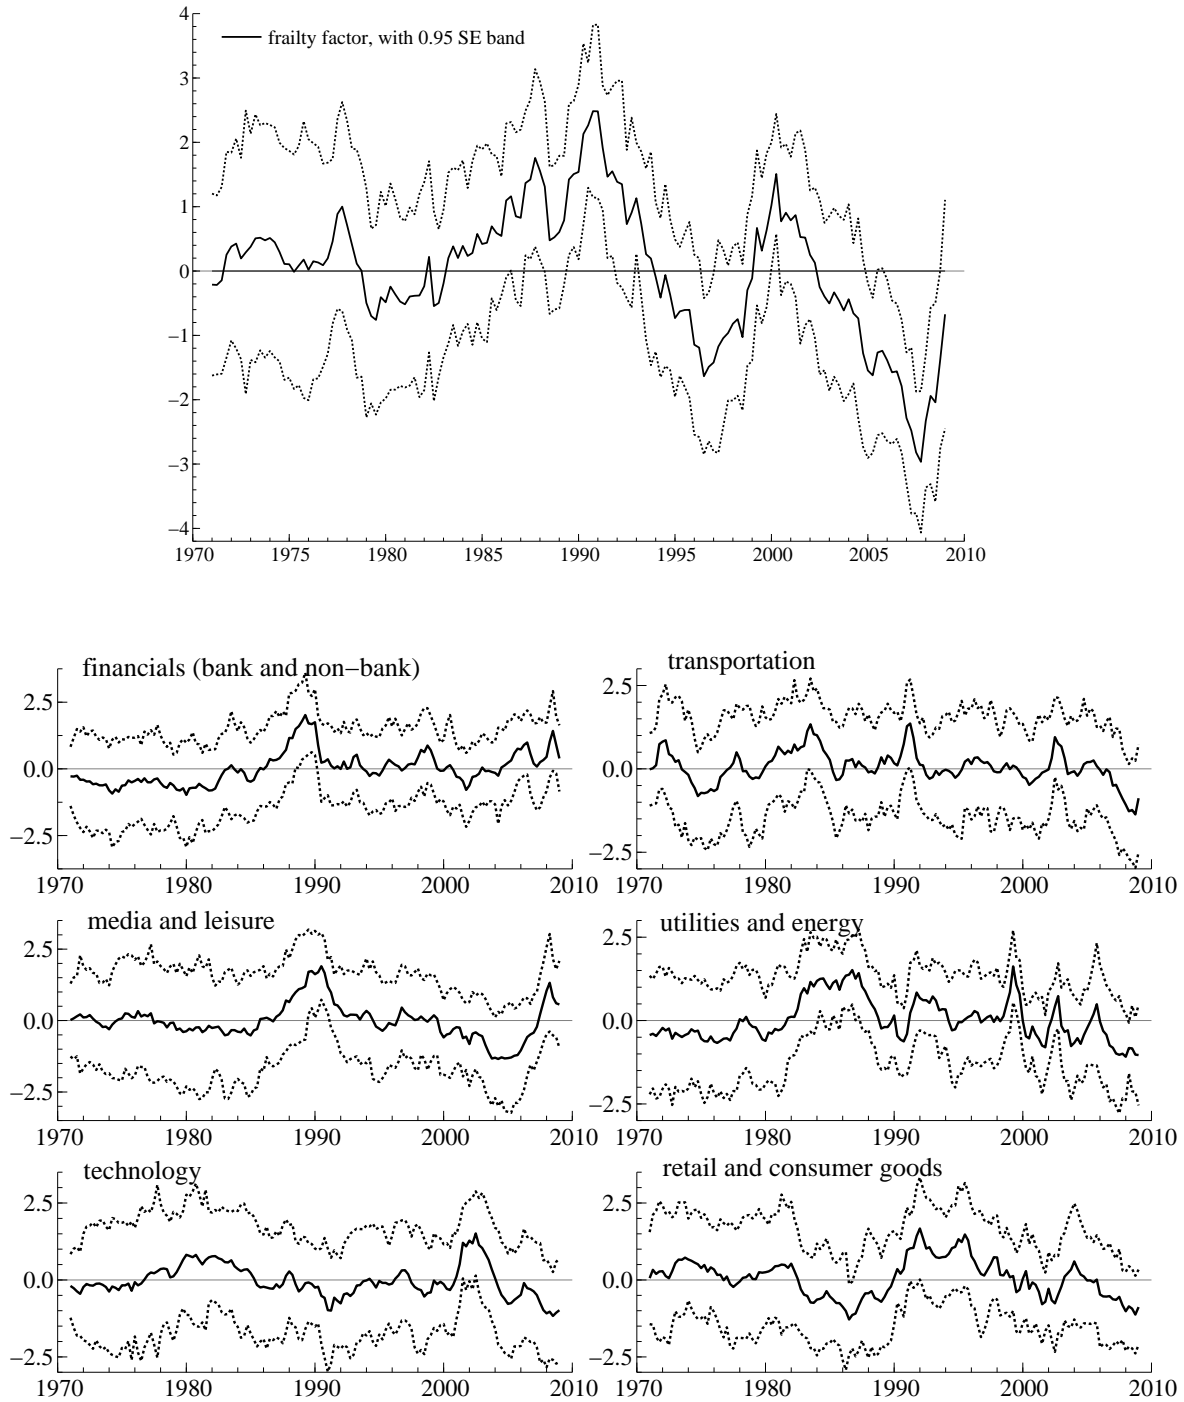



Figure 6: Estimated default signals

The left and right panels plot the time variation in default signals θ_{jt} for firms rated investment grade and speculative grade, respectively. In each panel, the share of variation implied by the macro factors f_t^m is indicated.




the default clustering implied by the ‘observed’ macro-financial risk factors is sufficient. For the speculative grade group, however, frailty effects imply additional default clustering during the late 1980s (savings and loan crisis), and help capture the very low default rates for bad risks in the years leading up to the financial crisis. This demonstrates that frailty effects are more important for financially weaker, and thus more credit constrained firms. Investment grade firms load on frailty to a relatively low extent, see also the parameter estimates in Table 2.

Second, Figure 7 relates changes in frailty effects to changes in bank lending standards (BLS) from the U.S. Senior Loan Officers (SLO) survey on lending practises, see (www.federalreserve.gov/boarddocs/snloansurvey). We consider the net percentage of banks that reported a tightening of credit standards for new commercial and industrial loans to large and medium sized enterprises. Lending standards are considered an imperfect but mainly accurate measure of credit supply, as demand conditions may also play a role, see e.g. Lown and Morgan (2006) and Maddaloni and Peydro (2011). Figure 7 shows a positive correlation between past net changes in lending standards on new loans and changes in estimated frailty effects. The intuition is that providing credit to increasingly worse borrowers

Figure 7: Frailty and bank lending standards

The left panel compares changes in the frailty factor with the net percentage of U.S. financial firms that have reported a tightening of lending standards in the Senior Loan Officer's survey. The right panel reports the corresponding scatterplot. The respective correlation coefficient is $\rho = 0.32$; the slope of the regression line is statistically significant at a 1% significance level.

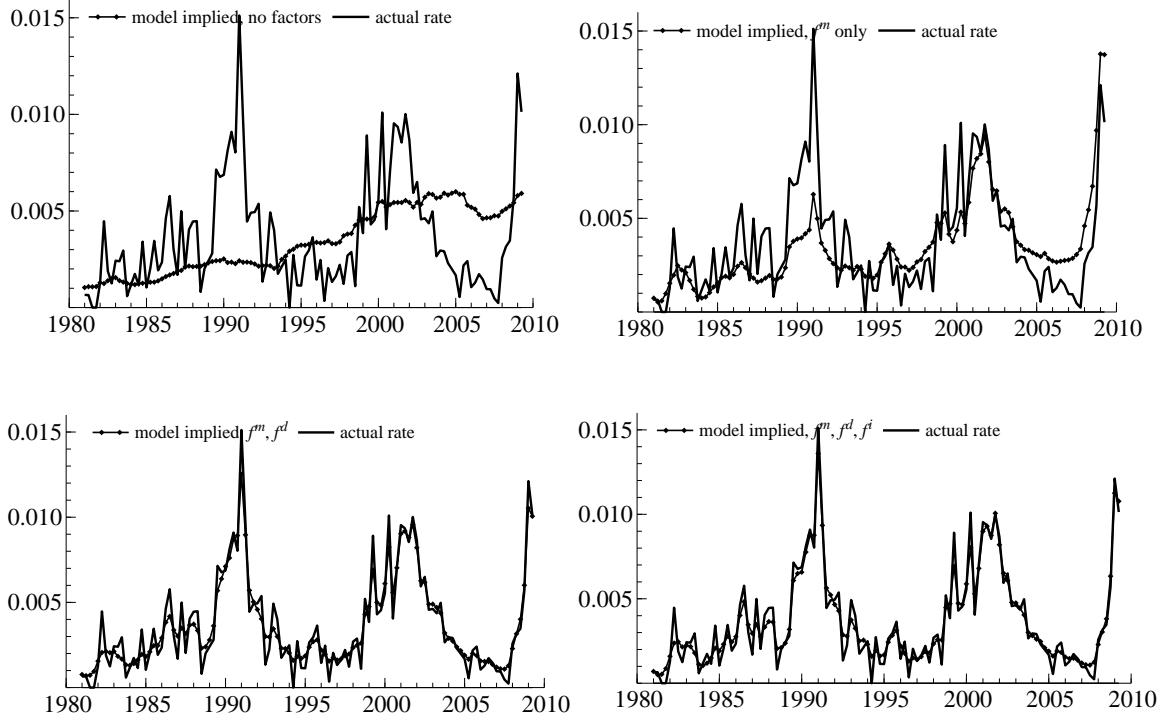


pushes default activity beyond what is implied by macro-financial covariates that are usually employed in credit risk modeling. The correlation is consistent with the notion frailty effects may capture specific omitted variables related to changes in credit supply, which are hard to measure empirically. Finally, frailty may also capture contagion effects across industry sectors. We refer to Azizpour et al. (2010) who develop a framework to disentangle contagion dependence from frailty.

Industry factors f_t^i capture deviations of industry-specific conditions from what is implied by purely common variation. Two alternative explanations come to mind. First, industry-specific dynamics may capture the industry-specific response of firms to common shocks. Second, intra-industry contagion through for example supply chain links may be an important source of dependence even after conditioning on common factors. Lando and Nielsen (2009) state that it is very difficult to find evidence for direct default contagion in the Moody's database. They base their conclusion on studying the qualitative summaries of individual firms' default histories. Further data analysis suggests that direct contagion plays no role, while some weaker evidence for indirect contagion through balance sheet covariates

Figure 8: Model fit to observed aggregate default rate

Each panel plots the observed quarterly default rate for all rated firms against the default rate implied by different model specifications. The models feature either (a) no factors, (b) only macro factors f^m , (c) macro factors and a frailty component f^m, f^d , and (d) all factors f^m, f^d, f^i , respectively.



exists. We come to a similar conclusion. For example, the default stress for technology firms in 2001-02 is clearly visible in the estimated industry-specific risk factor, but is most likely not due to contagion, but to the burst of the earlier asset bubble. Similarly, the 9/11 shock to the airline industry is visible as a brief spike in the transportation sector at the time, and difficult to interpret as contagion. As a result, the statistical and economic significance of factors f_t^i is more likely due to industry-specific heterogeneity rather than domino-style contagion dynamics. We accept, however, that no more rigorous distinction can be made based on the available data. We again refer to Azizpour, Giesecke, and Schwenkler (2010) for a first step to empirically distinguish common factors from contagious spillovers.

Figure 8 presents the model-implied economy-wide default rate against the aggregate

observed rates. We distinguish four specifications with (a) no factors, (b) f_t^m only, (c) f_t^m, f_t^d , and (d) all factors f_t^m, f_t^d, f_t^i . Based on these specifications, we assess the goodness of fit achieved at the aggregate level when adding latent factors. The static model fails to capture the observed default clustering around recession periods. The changes in the default rate for the static model are due to changes in the composition and quality of the rated universe. Such changes are captured by the rating and industry specific intercepts in the model. The upper-right panel indicates that the inclusion of macro variables helps to explain default rate variation. The latent frailty dynamics given by f_t^d , however, are clearly required for a good model fit. This holds both in low default periods such as 2002-2007, as well as in high default periods such as 1991. The bottom graphs of Figure 8 indicate that industry-specific developments cancel out in the cross-section to some extent and can thus be diversified. As a result, they may matter less from a (fully diversified) portfolio perspective.

3.2 Total default risk: a decomposition

We use the pseudo- R^2 measure as explained in Section 2.4 to assess which share of default rate volatility is captured by an increasing set of systematic risk factors. We are the first to do so in detail and detail the changing composition of systematic default risk over time.

Table 3 reports the estimated risk shares. By pooling over rating and industry groups, and by taking into account default and macroeconomic data for more than 35 years, we find that approximately 66% of a firm's total default risk is idiosyncratic, and 8.6% is industry related. The idiosyncratic risk (and to some extent the industry risk as well) can be eliminated in a large credit portfolio through diversification. The remaining share of risk, approximately 25%, does not average out in the cross section and is referred to as systematic risk. We find that for financially healthy firms (high ratings) the largest share of systematic default risk is due to the common exposure to macroeconomic and financial time series data. This common exposure can be regarded as the business cycle component. It constitutes approximately 58%

of systematic risk for firms rated investment grade, and 30–37% for firms rated speculative grade. The business cycle variation is not sufficient to account for all default rate variability in the data. Specifically, our results indicate that approximately 14% of total default risk, which is 41% of systematic risk, is due to an unobserved frailty factor. Frailty risk is low for investment grade firms (6%), but substantially larger for financially weaker firms (for 26% for Caa to 53% for B rated firms). Finally, approximately 9% of total default risk, or 25% of systematic risk, can be attributed to industry-specific developments.

Table 3 indicates how the estimated risk shares vary across rating and industry groups. The question whether firms rated investment grade have higher systematic risk than firms rated speculative grade is raised for instance by the Basel committee, see Basel Committee on Banking Supervision (2004). The Basel II framework imposes lower asset correlations for financially weaker firms, indicating lower systematic risk. Empirical studies employing a single latent factor tend to confirm this finding, see McNeil and Wendin (2007), and Koopman and Lucas (2008). In contrast to earlier studies, the last column of Table 3 indicates that speculative grade firms do not have less systematic risk than investment grade firms. This finding can be traced back to two sources. First, the frailty factor loads more heavily on speculative grade firms than investment grade firms. Second, some macro risk factors load on low rating groups also, see Table 2.

Figure 9 presents time series of estimated risk shares over a rolling window of eight quarters. These estimated risk shares vary considerably over time. While common variation with the business cycle explains approximately 11% of total variation on average, this share may be as high as 40%, for example in the years leading up to 2007. Similarly, the frailty factor captures a higher share of systematic default risk before and during times of crisis such as 1990-1991 and 2006-2007. In the former case, positive values of the frailty factor imply higher default rates that go beyond those implied by macroeconomic data. In the latter case, the significantly negative values of the frailty factor during 2006-2007 imply lower default

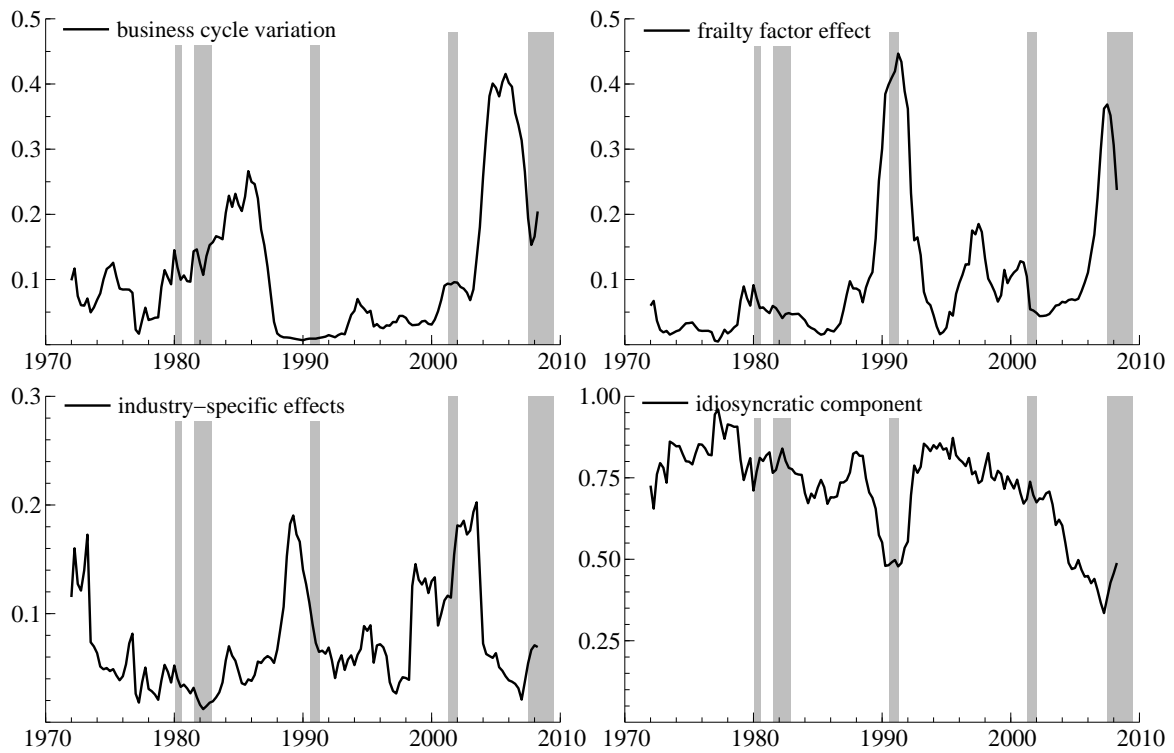
Table 3: A decomposition of total default risk

This table decomposes total (systematic and idiosyncratic) default risk into four unobserved constituents. We distinguish (i) common variation in defaults with observed macroeconomic and financial data, (ii) latent default-specific (frailty) risk, (iii) latent industry-sector dynamics, and (iv) non-systematic, and therefore diversifiable risk. The decomposition is based on our parameter estimates using data from 1971Q1 to 2009Q1.

Data	Business cycle f_t^c	Frailty risk f_t^d	Industry-level f_t^i	Idiosyncratic <i>distr.</i>
Pooled	11.4% (33.6%)	13.9% (40.9%)	8.6% (25.4%)	66.1%
Rating groups:				
Aaa-Baa	10.4% (58.0%)	1.1% (6.3%)	6.4% (35.7%)	82.1%
Ba	7.1% (34.0%)	7.5% (36.0%)	6.2% (30.0%)	79.2%
B	12.5% (30.0%)	22.3% (53.2%)	7.0% (16.8%)	58.2%
Caa-C	12.3% (36.7%)	8.9% (26.5%)	12.3% (36.8%)	66.5%
Industry sectors:				
Bank	5.4%	11.9%	18.8%	63.8%
Financial non-Bank	5.0%	5.3%	9.2%	80.5%
Transportation	7.4%	13.7%	18.8%	60.1%
Media	10.6%	19.9%	8.8%	60.8%
Leisure	15.7%	11.1%	2.6%	70.7%
Utilities	1.1%	4.9%	10.7%	83.3%
Energy	24.0%	8.7%	18.0%	49.3%
Industrial	16.3%	23.1%	-	60.7%
Technology	17.2%	11.0%	12.5%	59.3%
Retail	6.7%	9.6%	10.4%	73.2%
Consumer Goods	4.6%	18.4%	1.3%	75.7%
Misc	4.5%	13.2%	1.4%	80.9%

Figure 9: Time variation in risk shares

We present risk shares estimated over a rolling window of eight quarters from 1971Q1 to 2009Q1. Shaded areas correspond to recession periods as dated by the NBER.



rates than expected from macroeconomic data only. High absolute values of the frailty factor imply times when systematic default risk diverges from business cycle developments as represented by the common factors. Industry specific effects have been important mostly during the late 1980s and 2001-02. These are periods when banking specific risk and the burst of the technology bubble are captured through industry-specific factors, respectively.

The bottom right graph of Figure 9 presents the share of idiosyncratic risk over time. We observe a gradual decrease in idiosyncratic risk building up to the 2007-2009 crisis. Defaults become more systematic between 2001 and 2007 due to both macro and frailty effects. Negative values of the frailty risk factor during these years indicate that default rates were ‘systematically lower’ than what would be expected from macroeconomic developments. The eight-quarter rolling R^2 for the macro factors decreases by a factor 2 from 40% to 20% over 2005Q1-2007Q4, further suggesting that the default cycle has decoupled from macro developments in the years leading up to the crisis. Given the rolling window approach, the instantaneous effect may be even higher. The correction of this phenomenon over the financial crisis is also visible in the graphs. Again, this underlines the need for default risk models that include other risk factors above and beyond standard observed macroeconomic and financial time series. Such factors pick up rapid changes in the credit climate that might not be captured sufficiently well by observed risk factors. We address the economic impact of frailty and industry factors in Section 4.

4 Implications for risk management

Many default risk models that are employed in day-to-day risk management rely on the assumption of conditionally independent defaults, or doubly stochastic default times, see Das et al. (2007). At the same time, most models do not allow for unobserved risk factors and intra-industry dynamics to capture excess default clustering. We have reported in Section 3.2

that frailty and industry factors often account for more than half of systematic default risk. In this section we explore the consequences for portfolio credit risk when frailty and industry factors are not accounted for in explaining default variation. This is of key importance for internal risk assessment as well as external (macro-prudential) supervision: If shared effects are missing, standard portfolio credit risk models may tend to be wrong all at the same time.

4.1 The frailty factor

The frailty factor captures a substantial share of the common variation in disaggregated default rates at the industry and rating level, see Table 3. The presence of a frailty factor may increase default rate volatility compared to a model without latent risk dynamics. As a result it may shift probability mass of the portfolio credit loss distribution towards more extreme values. This would increase the capital buffers prescribed by the model. To explore this issue we conduct the following stylized credit risk experiment.

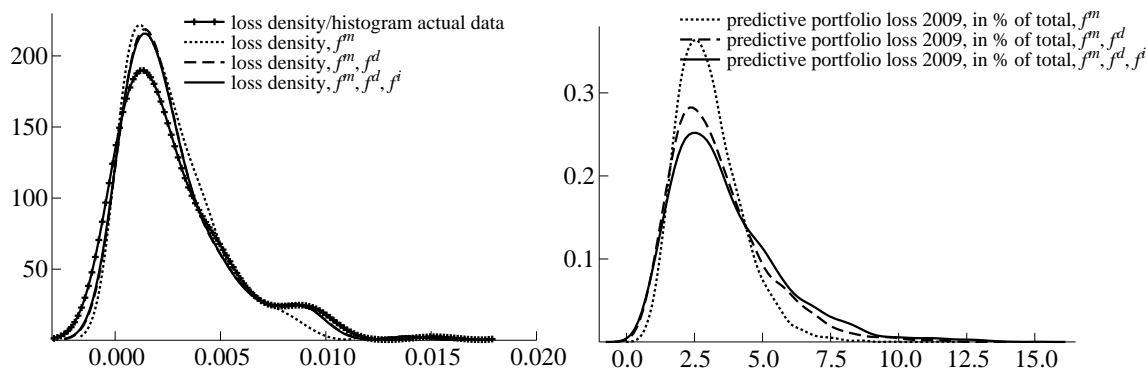
We consider a portfolio of short-term (rolling) loans to all Moody's rated U.S. firms. Loans are extended at the beginning of each quarter during 1981Q1 and 2008Q4 at no interest. A non-defaulting loan is re-extended after three months. The loan exposure to each firm at time t is given by the inverse of the total number of firms at that time, that is $(\sum_j k_{jt})^{-1}$. This implies that the total credit portfolio value is 1\$ at all times. For simplicity, we assume a stressed loss-given-default of 80%.

This example portfolio is stylized in many regards. Nevertheless, it allows us to investigate the importance of macroeconomic, frailty, and industry-specific dynamics for the risk measurement of a diversified loan or bond portfolio.

It is straightforward to simulate the portfolio credit loss distribution and associated risk measures for arbitrary credit portfolios in our setting. First, the exposures k_{jt} are chosen to correspond to the portfolio exposures. Second, one uses the methods introduced in the Appendix to simulate the current position of the latent systematic risk factors. Third, one

Figure 10: Real vs. model-implied credit portfolio loss distribution

We present distribution plots for a credit portfolio with uniform loan exposures to Moody's rated firms. The left panel presents the unconditional loss distribution as implied by historical quarterly defaults and firm counts in the database. The horizontal axis measures quarterly loan losses as a fraction of portfolio value. The left panel also presents the unconditional loss densities as implied by models with macro factors f_t^m , macro factors and a frailty component f_t^m, f_t^d , and all factors f_t^m, f_t^d, f_t^i , respectively. Positive probabilities of a negative portfolio loss are due to the (Gaussian Kernel) smoothing of the histogram. The right panel presents three simulated predictive portfolio loss densities for the year 2009, conditional on macro and default data until end of 2008, for different risk factor specifications. The horizontal axis measures annual losses as fractions of portfolio value.



can use (2) directly to simulate future risk factor realizations. Finally, conditional on the risk factor path, the defaults can be simulated by combining (3) and (4). Term structures of default rates can easily be obtained by combining model-implied quarterly probabilities over time. Out of sample forecasting exercises lie outside the scope of this paper. However, Azizpour et al. (2010) and Koopman et al. (2011) argue that credit risk models with both observed and unobserved risk factors forecast well.

The left panel in Figure 10 contains the credit portfolio loss distribution implied by actual historical default data. Since loss-given-default is held constant at 80%, this loss density is a horizontally scaled version of historically observed losses based on the default of then active firms. This distribution can be compared with the (unconditional) loss distribution implied by three different specifications of our model from Section 2. Portfolio loss densities for actual loan portfolios are known to be skewed to the right and leptokurtic, see e.g. McNeil, Frey, and Embrechts (2005, Chapter 8). Flat segments or bi-modality may arise due to the discontinuity in recovered principals in case of default. These qualitative features are

confirmed in the top panel of Figure 10.

By comparing the unconditional loss distributions in the left panel of Figure 10, we find that the common variation obtained from macroeconomic data is in general not sufficient to reproduce the thick right-hand tail implied by actual default data. In particular the shape of the upper tail of the empirical distribution is not well reproduced if only macro factors are used. The additional frailty and industry factors shift some of the probability mass into the right tail. The loss distributions implied by these models are closer to the actual distribution. The full model is able to reproduce the distributional characteristics of default rates, such as the positive skewness, excess kurtosis, and an irregular shape in the upper tail.

The right-hand panel of Figure 9 plots the simulated predictive credit portfolio loss densities for the year 2009, conditional on data until the end of 2008, as implied by different model specifications. Similarly to the unconditional case, the frailty factor shifts probability mass from the center of the distribution into the upper tail. Simulated risk measures are higher as a result. For the plotted densities, the simulated 99th percentile shifts out from about 6.24% to 8.34% of total portfolio value, which is an increase by more than 33%. Predicted annual mean losses are comparable at 2.96% and 2.71%, respectively. This demonstrates that frailty effects are economically important in addition to being statistically significant.


4.2 Industry specific risk dynamics

Section 3.2 shows that industry-specific variation accounts for about 25% of systematic default rate variation at the rating and industry level. Industry-specific risk factors capture the differential impact of macroeconomic developments on a given sector.

A specific case illustrates how macro, frailty, and industry-specific dynamics combine to capture industry-level variation in default rates. Figure 11 presents the observed quarterly default rate for the financial sector subsample of the entire Moody's data base. The rate is computed as the percentage of financial sector defaults over the total number of firms rated

Figure 11: Quarterly time-varying default intensities for financial firms

We present smoothed estimates of quarterly time-varying default rates for the financial sector. We distinguish a model with (i) common variation with macro data only, (ii) macro factors and a frailty component, and (iii) macro factors, frailty component, and industry-specific factors, respectively. The model-implied quarterly rates are plotted against the observed default fractions for financial firms.



in the financial industry. We distinguish three model specifications for the common variation, with macroeconomic factors only, with macro and frailty factors, and with macro, frailty, and industry-specific factors. Clearly, government intervention and bailouts have an effect on financial sector defaults, in particular at the end of our sample. Also, the objectives of a policy maker may be different from those of an investor in corporate bonds, which means that alternative definitions of default may be appropriate depending on the aim of the study. We proceed with Moody's definition of default, and focus on the relative importance of each set of factors over the last 30 years.

Figure 11 demonstrates that systematic variation of defaults due to shared exposure to common macro-financial covariates captures a substantial share of the overall time-variation in financial sector default rates. Also, the frailty factor is of key importance: It captures the overall default activity that is higher before and during the 1991 and 2001 recessions, and substantially lower in the years 2005-2007, see Figure 5. Finally, the industry-specific factor for financials, as plotted in the second panel of Figure 5, captures the additional sector-specific stress during the banking crisis periods of 1986-1990 and to some degree in 2008. It also adjusts the default rate (downwards) to the observed lower rates during the

2001 recession.

We conclude that industry factors are an important source of variation for defaults at the industry level. The bottom graphs of Figure 10 indicate that industry-specific developments may cancel to some extent, at least in a large loan portfolio that is also diversified across industries. If a portfolio is less well diversified, however, and exhibits clear industry concentrations, industry-specific effects may be a dominant cause for additional default clustering.

5 Conclusion

We have presented a new decomposition framework for systematic default risk. By means of a dynamic factor analysis, we can measure the contribution of macro, frailty, and industry-specific risk factors to overall default rate volatility. In our study of defaults for U.S. firms, we found that approximately one third of default rate volatility at the industry and rating level is systematic. The systematic default rate volatility can be further decomposed into macro and frailty driven. The part due to dependence on common macroeconomic conditions and financial activity ranges from about 30% for subinvestment up to 60% for investment grade companies. The remaining share of systematic credit risk is captured by frailty and industry risk. These findings suggest that credit risk management at the portfolio level should account for observed macro drivers of credit risk as well as for unobserved risk factors. In particular, standard portfolio credit risk models that account for macroeconomic dependence only leave out a substantial part of systematic credit risk.

We have given further empirical evidence that the composition of systematic risk varies over time. In particular, we observe a gradual build-up of systematic risk over the period 2002-2007. Such patterns can be used as early warning signals for financial institutions and supervising agencies. If the degree of systematic comovement between credits exposures

increases through time, the fragility of the financial system may increase and prompt for an adequate (re)action. Interestingly, the frailty component appears more important for lower grade companies in periods of stress, which is precisely the time when such companies might be expected to experience more difficulties in rolling over debt. In addition, changes in our estimated frailty factor appear positively correlated with tightening lending standards. This suggests that frailty factors may capture changes in economic conditions which are hard to quantify, but impact the quality and default experience of bank portfolios in an economically significant way.

Our results have a clear bearing for risk management at financial institutions. When conducting risk analysis at the portfolio level, the frailty and industry components cannot be discarded. This is confirmed in a risk management experiment using a stylized loan portfolio. The extreme tail clustering in defaults cannot be captured using macro variables alone. Additional sources of default volatility such as frailty and contagion need to be identified in order to capture the patterns in default data over time.

References

- Alessi, L., M. Barigozzi, and M. Capasso (2008). A robust criterion for determining the number of static factors in approximate factor models. *ECB working paper* (903), 1–27.
- Azizpour, S., K. Giesecke, and G. Schwenkler (2010). Exploring the sources of default clustering. *Stanford University working paper series*.
- Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica* 70(1), 191–221.
- Basel Committee on Banking Supervision (2004). Basel II: International Convergence of Capital Measurement and Capital Standards: a Revised Framework. *Bank of International Settlements Basel Report 107*.

- Cameron, A. C. and F. A. G. Windmeijer (1997). An R squared measure of goodness of fit for some common nonlinear regression models. *Journal of Econometrics* 77, 329–342.
- CEBS (2010). Committee of European Banking Supervisors: Summary Report from 23 July 2010. *available at www.eba.europa.eu*.
- Das, S., D. Duffie, N. Kapadia, and L. Saita (2007). Common Failings: How Corporate Defaults Are Correlated. *Journal of Finance* 62(1), 93–117(25).
- Duffie, D., A. Eckner, G. Horel, and L. Saita (2009). Frailty Correlated Default. *Journal of Finance* 64(5), 2089–2123.
- Duffie, D., L. Saita, and K. Wang (2007). Multi-Period Corporate Default Prediction with Stochastic Covariates. *Journal of Financial Economics* 83(3), 635–665.
- Durbin, J. and S. J. Koopman (1997). Monte Carlo Maximum Likelihood estimation for non-Gaussian State Space Models. *Biometrika* 84(3), 669–684.
- Durbin, J. and S. J. Koopman (2001). *Time Series Analysis by State Space Methods*. Oxford: Oxford University Press.
- Durbin, J. and S. J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika* 89(3), 603–616.
- Figlewski, S., H. Frydman, and W. Liang (2008). Modeling the Effect of Macroeconomic Factors on Corporate Default and Credit Rating Transitions. *New York University Discussion Paper*.
- Frey, R. and A. J. McNeil (2002). VaR and Expected Shortfall in Portfolios of Dependent Credit Risks: Conceptual and Practical Insights. *Journal of Banking and Finance* 26(7), 1317–1334.
- Geweke, J. (1989). Bayesian inference in econometric models using Monte Carlo integration. *Econometrica* 57, 1317–39.
- Giesecke, K. (2004). Correlated default with incomplete information. *Journal of Banking and Finance* 28, 1521–1545.
- Giesecke, K. and S. Azizpour (2008). Self-exciting corporate defaults: contagion vs frailty. *Stanford University working paper series*.

- Jorion, P. and G. Zhang (2007). Good and bad credit contagion: evidence from credit default swaps. *Journal of Financial Economics* 84(3), 860–883.
- Kloek, T. and H. K. van Dijk (1978). Bayesian estimates of equation system parameters: an application of integration by Monte Carlo. *Econometrica* 46, 1–20.
- Koopman, S., R. Kräussl, A. Lucas, and A. Monteiro (2009). Credit cycles and macro fundamentals. *Journal of Empirical Finance* 16, 42–54.
- Koopman, S. J. and A. Lucas (2008). A Non-Gaussian Panel Time Series Model for Estimating and Decomposing Default Risk. *Journal of Business and Economic Statistics* 26 (4), 510–525.
- Koopman, S. J., A. Lucas, and A. Monteiro (2008). The Multi-Stage 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.
- Koopman, S. J., N. Shephard, and D. Creal (2009). Testing the assumptions behind importance sampling. *Journal of Econometrics* 149(1), 2–10.
- Lando, D. and M. S. Nielsen (2009). Correlation in corporate defaults: Contagion or conditional independence. *Copenhagen Business School, working paper*.
- Lang, L. and R. Stulz (1992). Contagion and competitive intra-industry effects of bankruptcy announcements. *Journal of Financial Economics* 32, 45–60.
- Lown, C. and D. P. Morgan (2006). The credit cycle and the business cycle: new findings using the the loan officer opinion survey. *Journal of Money, Credit and Banking* 38(6), 1575–97.
- Maddaloni, A. and J.-L. Peydro (2011). Bank risk taking, securitization, supervision and low interest rates: evidence from the euro area and the U.S. lending standards. *Review of Financial Studies*, forthcoming 1248, 1–53.
- 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.

- McNeil, A. J., R. Frey, and P. Embrechts (2005). *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press.
- Pesaran, H., T. Schuermann, B. Treutler, and S. Weiner (2006). Macroeconomic Dynamics and Credit Risk: A Global Perspective. *Journal of Money, Credit, and Banking* 38, No. 5, 1211–1261.
- Shephard, N. and M. K. Pitt (1997). Likelihood analysis of non-Gaussian measurement time series. *Biometrika* 84, 653–67.
- Stock, J. and M. Watson (2002). Macroeconomic Forecasting Using Diffusion Indexes. *Journal of Business and Economic Statistics* 20(2), 147–162.
- Tarullo, D. K. (2010). Equipping Financial Regulators with the Tools Necessary to Monitor Systemic Risk. *statement before the Subcommittee on Security and International Trade and Finance, Committee on Banking, Housing, and Urban Affairs, U.S. Senate, 12 February 2010*.