Macro, Industry and Frailty Effects in Defaults:
The 2008 Credit Crisis in Perspective

Siem Jan Koopman
André Lucas
Bernd Schwaab

VU University Amsterdam, and Tinbergen Institute.

Macro, frailty, and contagion effects in defaults: lessons from the 2008 credit crisis *

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

August 24, 2010

Macro, frailty, and contagion effects in defaults: lessons from the 2008 credit crisis

Abstract

In the aftermath of the financial crisis, banks in the U.S. and Europe have been subjected to a sequence of stress tests to measure system stability. Such tests are formulated in terms of adverse economic scenarios rather than in terms of systematic default rate increases. This suggests that macroeconomic conditions fully capture default stress. However, two additional explanations can be found in the literature for the occurrence of default clusters: autonomous default rate dynamics, also known as frailty, and industry-specific effects including contagion. We develop a new methodological framework to disentangle, quantify, and test these three competing explanations. Using U.S. default data we find that observed macro and financial market factors account for only 30–60% of systematic default risk. Consequently, stress-testing frameworks that only control for observed macro conditions leave out a substantial share of systematic risk. The components not related to business-cycle dynamics (frailty) are particularly relevant before and during times of financial market turbulence. For example, we find clear systemic risk build-up over the period preceding the 2008 credit crisis.

1 Introduction

During and following the 2008 financial crisis, banks in the U.S. and Europe have been subjected to a range of stress tests in an effort to assess the safety and solidity of the banking system. Such tests typically involve extreme scenarios for the main risk drivers of banks, such as economic growth figures, default rates, interest and exchange rates, and liquidity levels. When focusing on the banking book, the dominant risk factor is the systematic variation in the default rate. Systematic increases in default rates may lead to clusters of defaults. 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. Typically, however, stress tests are not formulated directly in terms of higher default rates, but rather in terms of economic scenarios. For example, the stress scenario in the 2010 European stress testing exercise involved a double-dip recession, with GDP growth dropping around 1% and 2% below the baseline prediction in 2010 and 2011, respectively. Additional stress in terms of rising unemployment rates, rising sovereign yields, and falling house prices was added to this. Based on these stressed economic conditions, banks were left to translate the economic scenarios into default rate scenarios, credit losses, and finally capital implications.

It is 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 macro economic conditions, however, 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.

Empirically understanding the origins of systematic default rate variation is thus of prime importance. If in addition to macroeconomic conditions there are additional factors that impact default clustering, current stress testing methodologies might be misguided and provide a misleading picture of the safety of the banking system. A different framework should then be designed that also accounts for the additional factors of default rate volatility. To assess the severity of this problem, we need a qualification and a quantification of the relative magnitude of the additional factors compared to standard macroeconomic risk drivers. A proper methodology for this, however, is currently lacking.

1

In this paper, we fill this gap by empirically disentangling the relative contribution of three different explanations for default rate clustering using U.S. corporate default data of Moody's. We develop a new methodological framework to study simultaneous factor structures for macroeconomic and financial data as well as discrete default counts. The framework by itself also has applications beyond the current context. Using the new framework, we quantify for the first time the relative contribution of the different origins of systematic default rate variation. We do so by constructing simple R^2-type measures based on our factor models and by measuring the increases in R^2 between the different empirical models.

The two additional explanations for default rate clustering that we consider in this paper besides the exposure to shared macroeconomic risk factors, are known as frailty-correlated defaults and contagion, respectively. In a frailty model, the additional variation in default intensities is captured by an unobserved dynamic (frailty) 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). The frailty factor captures default clustering above and beyond what can be explained by macroeconomic variables and firm-specific information. The unobserved component can pick up the 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). Contagion effects may dominate potentially off-setting competitive effects at the intra-industry level, see e.g. Lang and Stulz (1992) and Jorion and Zhang (2007). Lando and Nielsen (2008) screen hundreds of default histories for evidence of direct default contagion. The examples suggest that contagion is mainly an intra-industry effect. As a result, contagion may explain default dependence at the industry level beyond that induced by macro and frailty factors.

Lando and Nielsen (2008) 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.

Our research indicates 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

2

of credit portfolio risk. 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, 30% 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 roughly 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. 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.

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 management 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

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

for $t = 1, \ldots, T$, where the first J elements of y_t are default counts. We count defaults for

different ratings and industries. As a consequence, the first J elements of y_t are discrete-valued. The remaining N elements of y_t contain macro and financial variables which take continuous values. We assume that both the default counts and the macro and financial time series data are driven by 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.

In our study, we distinguish macro, frailty, and industry (or contagion) factors. The 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 data and default counts. Therefore, factors f_t^m are common to all data. Frailty factors f_t^d are default-specific, i.e., common to default data $(y_{1t}, \ldots, y_{Jt})$ and independent of observed macroeconomic and financial data by construction. By not allowing the frailty factors to impact the macro series y_{jt} for $j = J+1, \ldots, J+N$, we effectively restrict f_t^d to pick up any default clustering above and beyond that is implied by macroeconomic and financial factors f_t^m. The third set of factors f_t^i considered in this paper affects firms in the same industry. Such factors may arise as a result of default contagion through up- and downstream business links. Alternatively, they may be interpreted as industry-specific frailty factors. Disentangling these two interpretations is empirically impossible unless detailed information at the firm-level is available on firm interlinkages at the trade and institutional level. Such data are not available for our current analysis.

We gather all factors into the vector $f_t' = (f_t^{m\prime}, f_t^{d\prime}, f_t^{i\prime})$. Note that we only observe the default counts and macro variables y_t. The factors f_t themselves are latent and thus unobserved. We assume the following simple autoregressive dynamics for the latent factors,

$$f_t = \Phi f_{t-1} + \eta_t, \qquad t = 1, 2, \ldots, \tag{2}$$

with the coefficient matrix Φ diagonal and with $\eta_t \sim \mathrm{N}(0, \Sigma_\eta)$. More complex dynamics than (2) can be considered as well. The autoregressive structure allows the components of f_t to be sticky. For example, it allows the macroeconomic factors f_t^m to evolve slowly over time and capture the business cycle component in both macro and default data. Similarly, the credit climate and industry default conditions can be captured by persistent processes for f_t^d and f_t^i, such that they can capture the clustering of high-default years. To complete the specification of the factor process, we specify the initial condition $f_1 \sim \mathrm{N}(0, \Sigma_0)$. We assume stationarity of the factor dynamics by insisting that all m eigenvalues of Φ lie inside the unit circle. The $m \times 1$ disturbance vectors η_t are serially uncorrelated.

To combine the normally and non-normally distributed elements in y_t, we adopt our mixed measurement approach. The MiMe DFM is based on the standard factor model

assumption: conditional on the factors f_t, the measurements in y_t are independent. In our specific case, we assume that conditional on f_t, the first J elements of y_t have a binomial distribution with parameters k_{jt} and π_{jt}, for $j = 1, \ldots, J$. Here, k_{jt} denotes the number of firms in a specific rating and industry bucket j at time t and π_{jt} denotes the probability of default conditional on f_t. For more details on the conditionally binomial model see e.g. McNeil, Frey, and Embrechts (2005, Chapter 9). Frey and McNeil (2002) show that all available industry credit risk models, i.e. Creditmetrics, KMV, CreditRisk+, can be presented as conditional binomial models. The remaining N elements of y_t follow conditional on f_t a normal distribution with mean μ_{jt} and variance σ_{jt}^2 for $j = J + 1, \ldots, J + N$.

2.1 The mixed measurement dynamic factor model

Both the binomial and the normal distribution are members of the exponential family of distributions. In this paper, we formulate the MiMe DFM for random variables from the exponential family. The model can easily be extended to handle distributions outside this class. The estimation methodology presented in this paper applies to the general case as well.

The link between the factors f_t and the observations y_t relies on time-varying location parameters, such as the default probability π_{jt} for default data and the mean μ_{jt} for Gaussian data. In general, let each variable y_{jt} follow the distribution

$$y_{jt}|\mathcal{F}_t \sim p_j(y_{jt}|\mathcal{F}_t; \psi), \tag{3}$$

where $\mathcal{F}_t = \{f_t, f_{t-1}, \ldots\}$ and ψ is a vector of fixed and unknown parameters that include, for example, the elements of Φ and Σ_η in (2). The index j of the density $p_j(\cdot)$ indicates that the type of measurement y_{jt} (discrete versus continuous) may vary across j. We assume that the information from past factors $\mathcal{F}_t$ impacts the distribution of y_{jt} through an unobserved signal θ_{jt}. For example, for the normal distribution, θ_{jt} equals the mean, while for the binomial case θ_{jt} is the log-odds ratio, $\log(\pi_{jt}/(1 - \pi_{jt}))$. For exponential family data, θ_{jt} is the so-called canonical parameter, see Appendix A1. We assume that the signal θ_{jt} is a linear function of unobserved factors, f_t, such that

$$\theta_{jt} = \alpha_j + \lambda_j' f_t, \tag{4}$$

with α_j an unknown constant, and λ_j an $m \times 1$ loading vector with unknown coefficients. It is conceptually straightforward to let θ_{jt} also depend on past values of the factors f_t. We emphasize that y_t may depend linearly as well as non-linearly on the common factors f_t.

As the key question in this paper concerns the relative contributions of macro, frailty, and contagion (or industry) risk to general default risk, we introduce further restrictions on the general form of (4). In particular, we specify the signals by

$$\theta_{jt} = \lambda_{0j} + \beta_j' f_t^m + \gamma_j' f_t^d + \delta_j' f_t^i, \quad \text{for } j = 1, \dots, J, \tag{5}$$

$$\theta_{jt} = \lambda_{0j} + \beta_j' f_t^m, \quad \text{for } j = J+1, \dots, J+N. \tag{6}$$

The signal specification in (6) implies that the means of the macro variables depend linearly on the macro factors f_t^m. The components of f_t^m capture general developments in business cycle activity, lending conditions, financial markets, etc. The log-odds ratios in (5) partly depend on macro factors, but also depend on frailty risk f_t^d and industry f_t^i factors. The specification of the signals in (5) and (6) is key to our empirical analysis where we focus on studying whether macro dynamics explain all systematic default rate variation, or whether and to which extend frailty and industry factors are also important.

For model identification, we impose the restriction $\Sigma = I - \Phi\Phi'$. This implies that the factor processes in (2) have an autoregressive structure with unconditional unit variance. It also implies that factor loadings in β_j, γ_j, and δ_j can be interpreted as factor standard deviations (volatilities) for firms of type $j = 1, \dots, J$.

As mentioned, all model parameters that need to be estimated are collected in a parameter vector ψ. This includes the factor loadings β_j, γ_j, δ_j, but also the coefficients in the autoregressive matrix Φ in (2). We aim to estimate ψ by maximum likelihood. For this purpose, we numerically maximize the likelihood function as given by

$$p(y; \psi) = \int p(y, f; \psi) \mathrm{d}f = \int p(y|f; \psi) p(f; \psi) \mathrm{d}f, \tag{7}$$

where $p(y, f; \psi)$ is the joint density of the observation vector $y' = (y_1', \dots, y_T')$ and the factors $f' = (f_1', \dots, f_T')$. The integral in (7) is not known analytically, and we therefore rely on numerical methods. The likelihood function (7) can be evaluated efficiently via Monte Carlo integration and using the method of importance sampling, see Durbin and Koopman (2001). Maximizing the Monte Carlo estimate of the likelihood function is feasible using standard computers. Once maximum likelihood estimates of ψ are obtained, (smoothed) estimates of the unobserved macro, frailty, and industry factors f_t and their standard errors can be obtained using the same Monte Carlo methods. This methodology has a number of interesting features in the current setting, but we defer all details on the estimation procedure to the Appendix.

2.2 Decomposition of non-Gaussian variation

Once the model parameters and risk factors are estimated, we need to assess which share of variation in default data is captured by the different latent factors. Obviously, this cannot be achieved by a standard R^2 measure. We therefore adopt a pseudo-R^2 measure which is similar to those discussed in Cameron and Windmeijer (1997). The pseudo-R^2 measure is based on a distance measure between two distributions. For the normal linear regression model, the pseudo-R^2 reduces to the familiar R^2 from regression.

Our distance measure for the pseudo-R^2 is the Kullback-Leibler (KL) divergence, which is defined as

$$KL(\theta_1, \theta_2) = 2 \int \left[\log p_{\theta_1}(y) - \log p_{\theta_2}(y)\right] p_{\theta_1}(y) \mathrm{d}y. \tag{8}$$

The KL divergence measures the average distance between the two log-densities $\log p_{\theta_1}$ and $\log p_{\theta_2}$, which are completely specified by parameter vectors θ_1 and θ_2, respectively. We are particularly interested in the pseudo-R^2 of the default equations of the model to measure the size and composition of systematic default risk. Therefore, in our current setting $p_\theta(y)$ is the binomial distribution for each rating-industry combination, while θ denotes the time series of corresponding (estimated) log-odds for that combination. The differences in log-odds are due to the use of different models.

Figure 1 illustrates the idea of assessing the contribution of common factors to default risk in more detail. We distinguish several alternative model specifications indicated by M^{na}, M^m, M^{md}, and M^{mdi}. These models contain an increasing collection of latent factors. Model M^{na} does not contain any factors, while models M^m, M^{md}, and M^{mdi} cumulatively add the macro, frailty, and industry factors, respectively. Model M^{max} provides the maximum possible fit by considering a model with a separate dummy variable for each observation. Thus, the model contains as many parameters as observations. While useless for practical purposes, the unrestricted model provides a natural benchmark for what is the maximum possible fit to the data.

The constructed log-odds can be substituted in (8) to decompose systematic credit risk. We consider the improvements in fit when moving from M^{na} to M^m, M^{md}, M^{mdi}, and ultimately to M^{max}. The pseudo-R^2 is now defined as

$$R^2(\theta) = 1 - \frac{KL(\theta_{max}, \theta)}{KL(\theta_{max}, \theta_{na})}. \tag{9}$$

Note that (9) scales the KL improvements by the total distance between models M^{max} and M^{na}, that is $KL(\theta_{max}, \theta_{na})$. This allows us to interpret (9) as the proportional reduction in variation due to the inclusion of latent factors, see Cameron and Windmeijer (1997).

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 increasing sets of latent factors.
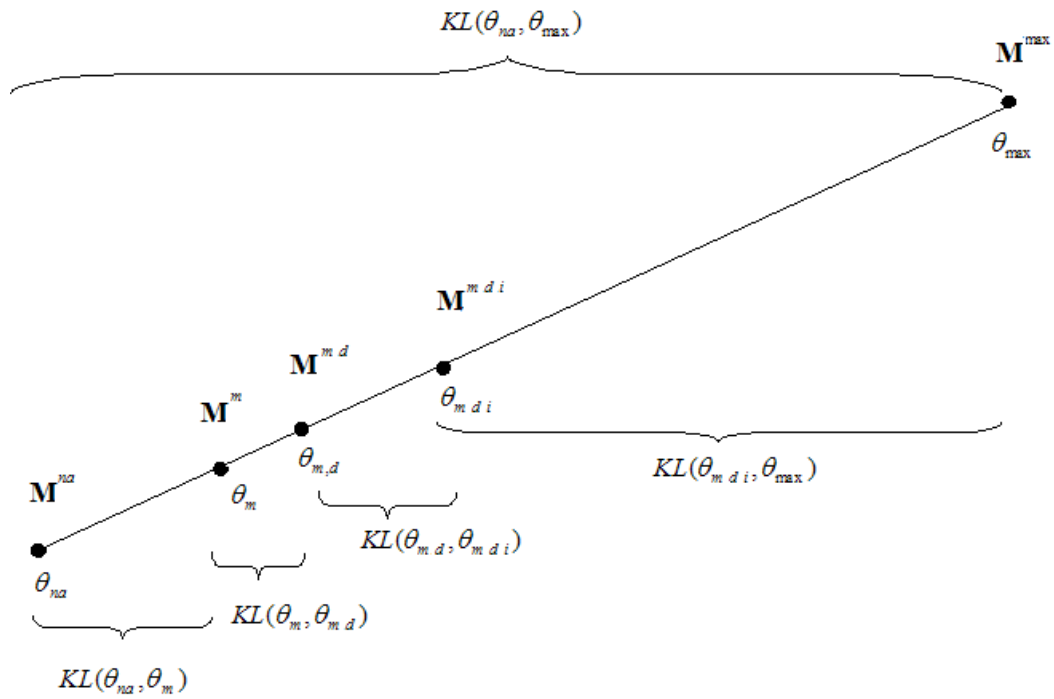
As mentioned earlier, for the standard linear regression model (9) reduces to the standard R^2. For the binary choice model, the McFadden pseudo-R^2 is obtained. Similar to the standard R^2, all values lie between zero and one. The relative contribution from each of our systematic credit risk factors can now be quantified by looking at the increase in pseudo-R^2 when moving from M^m via M^{md} to M^{mdi}. The remainder from M^{mdi} to M^{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. Whenever possible, we relate our findings to questions from the finance and credit risk literature that we perceive to be open issues. We distinguish seven industry groups (financials and insurance; transportation; media, hotels, and leisure; utilities and 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. In addition, ratings may be less noisy compared to raw balance sheet or equity market based data.

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. This data enters the analysis in the form of annual growth rates, see Figure 3 for time series plots.

Figure 2: Clustering in default data

The top graph plots (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 graph plots time series of default fractions y_{jt}/k_{jt} over time. We distinguish four broad rating groups, i.e., $Aaa - Baa$, Ba, B, and $Caa - C$, where each plot contains 12 time series of industry-specific default fractions.
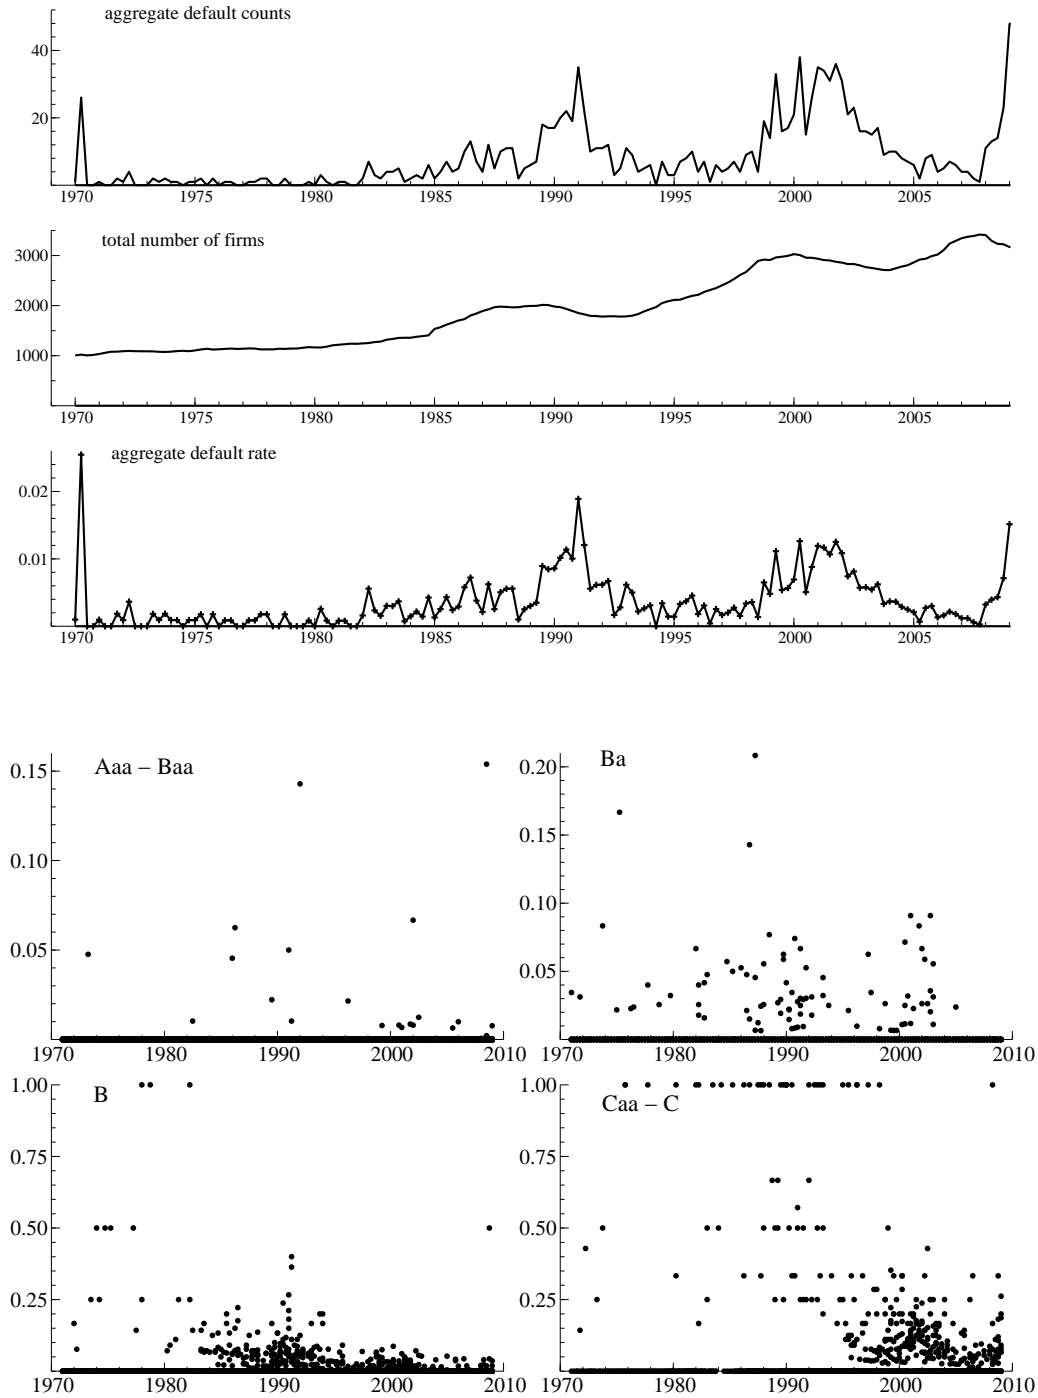
Table 1: Macroeconomic Time Series Data

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	Total no
(a) Macro indicators, and business cycle conditions	Industrial production index	
	Disposable personal income	
	ISM Manufacturing index	
	Uni Michigan consumer sentiment	5
	New housing permits	
(b) Labour market conditions	Civilian unemployment rate	
	Median duration of unemployment	
	Average weekly hours index	4
	Total non-farm payrolls	
(c) Monetary policy and financing conditions	Federal funds rate	
	Moody's seasoned Baa corporate bond yield	
	Mortgage rates, 30 year	
	10 year treasury rate, constant maturity	6
	Credit spread corporates over treasuries	
	Government bond term structure spread	
(d) Bank lending	Total Consumer Credit Outstanding	
	Total Real Estate Loans, all banks	2
(e) Cost of resources	PPI Fuels and related Energy	
	PPI Finished Goods	
	Trade-weighted U.S. dollar exchange rate	3
(f) Stock market returns	S&P 500 yearly returns	
	S&P 500 return volatility	2
		22

Figure 3: Macroeconomic and financial time series data

The graph contains times series plots of yearly growth rates in macroeconomic and financial data. For a listing of the data we refer to Table 1.
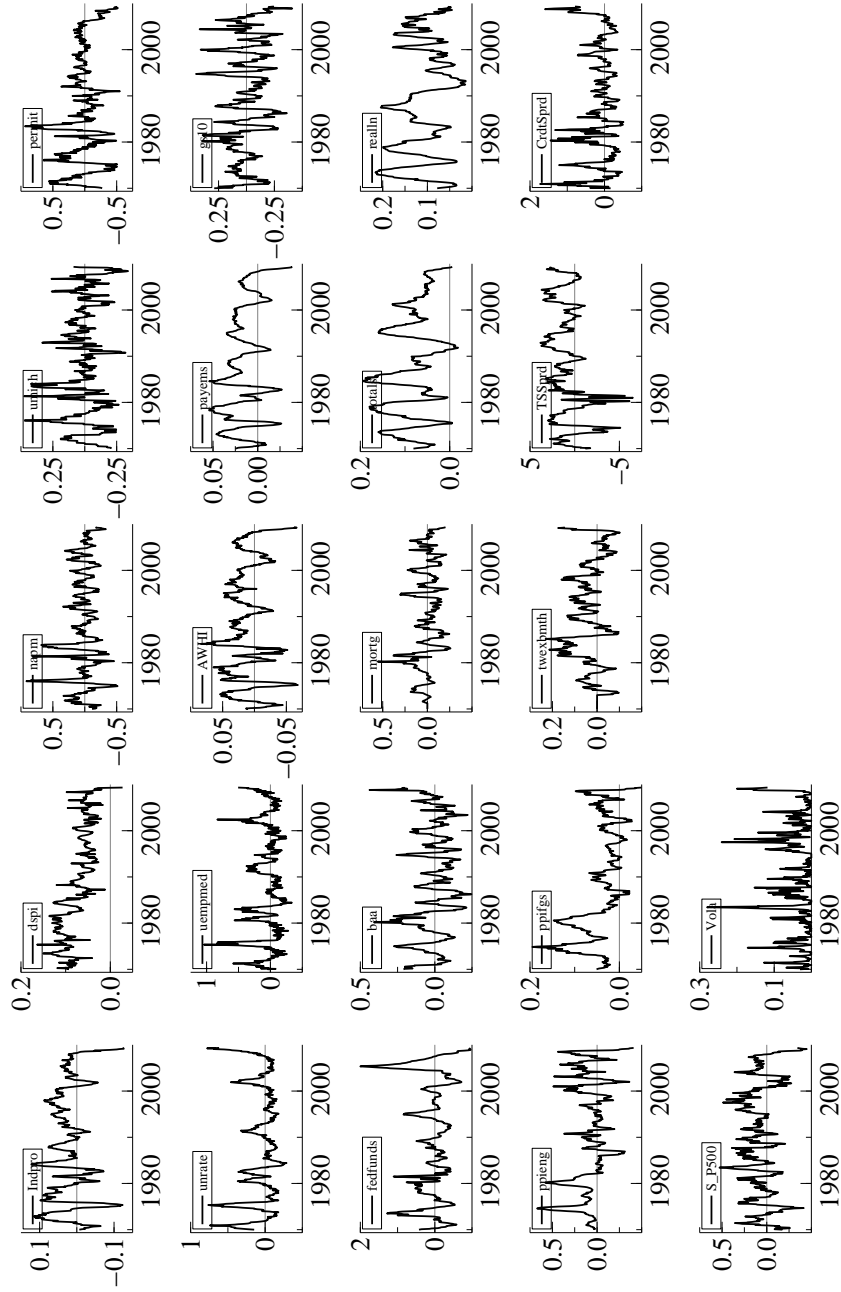
3.1 Major empirical results

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 $\lambda_{0,j} = \lambda_0 + \lambda_{1,r_j} + \lambda_{2,s_j}$, where the common intercept λ_0 is adjusted by specific coefficients indicating industry sector (s_j) and rating group (r_j), respectively, for $j = 1, \ldots, J$ 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 are 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 (or contagion) 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.

Figure 4 presents four estimated risk factors f_t^m as defined in (5) and (6). We graph 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 A2. 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.

Figure 5 presents the shares of variation in each macroeconomic time series that can be attributed to the common macroeconomic factors. The first two macroeconomic factors load mostly on labor market, production, and interest rate data. The last two factors displayed in the bottom panels of Figure 5 load mostly on survey sentiment data and changes in price level indicators. The macroeconomic factors capture 24.7%, 22.4%, 11.0%, and 8.0% of the total variation in the macro data panel, respectively (66.1% in total). The range of explained variation ranges from about 30% (S&P 500 index returns, fuel prices) to more than 90% (unemployment rate, average weekly hours index, total non-farm payrolls). 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 6 presents smoothed 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

Table 2: Parameter estimates, binomial part

We report parameter estimates associated with the binomial data. The coefficients in the first column combine to fixed effects according to $\lambda_{0,j} = \lambda_0 + \lambda_{1,r_j} + \lambda_{2,s_j}$, i.e., 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 loading coefficients 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 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	-2.51	7.61	$\beta_{1,IG}$	0.36	0.57	γ_{IG}	0.23	0.81
			$\beta_{1,Ba}$	0.23	0.56	γ_{Ba}	0.46	2.48
λ_{fin}	-0.23	1.03	$\beta_{1,B}$	0.25	0.64	γ_B	0.65	3.93
λ_{tra}	-0.06	0.24	$\beta_{1,C}$	0.15	0.66	γ_C	0.40	5.35
λ_{lei}	-0.21	0.86						
λ_{utl}	-0.68	2.13	$\beta_{2,IG}$	0.22	0.36	Loadings f_t^i		
λ_{tec}	-0.09	0.62	$\beta_{2,Ba}$	0.24	0.60	δ_{fin}	0.68	0.97
λ_{ret}	-0.36	1.73	$\beta_{2,B}$	0.06	0.16	δ_{tra}	0.64	0.78
			$\beta_{2,C}$	0.26	1.32	δ_{lei}	0.40	1.61
λ_{IG}	-7.10	15.70				δ_{utl}	0.99	2.01
λ_{BB}	-3.89	12.11	$\beta_{3,IG}$	0.74	1.62	δ_{tec}	0.40	1.62
λ_B	-2.12	9.59	$\beta_{3,Ba}$	0.44	1.54	δ_{ret}	0.40	1.67
			$\beta_{3,B}$	0.24	1.05			
			$\beta_{3,C}$	0.20	1.74			
			$\beta_{4,IG}$	0.67	0.93			
			$\beta_{4,Ba}$	0.10	0.23			
			$\beta_{4,B}$	-0.34	-0.73			
			$\beta_{4,C}$	-0.10	-0.44			

Figure 4: Smoothed Macroeconomic Risk Factors

This figure presents four estimated risk factors f_t^m as defined in (5) and (6). We plot the estimated conditional mean of the factors, along with approximate standard error bands at a 95% confidence level.
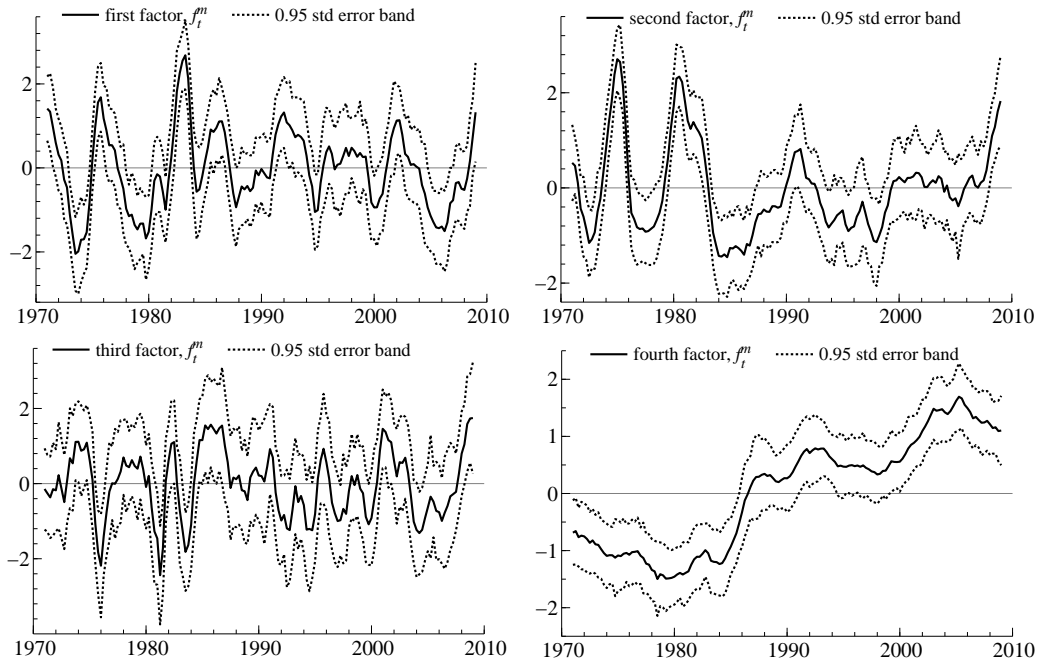
Figure 5: Shares of explained variation in macro and financial time series data
The figure indicates which share of variation in each time series listed in Table 1 can be attributed to each factor f^m. Factors f^m are common to the (continuous) macro and financial as well as the (discrete) default count data.
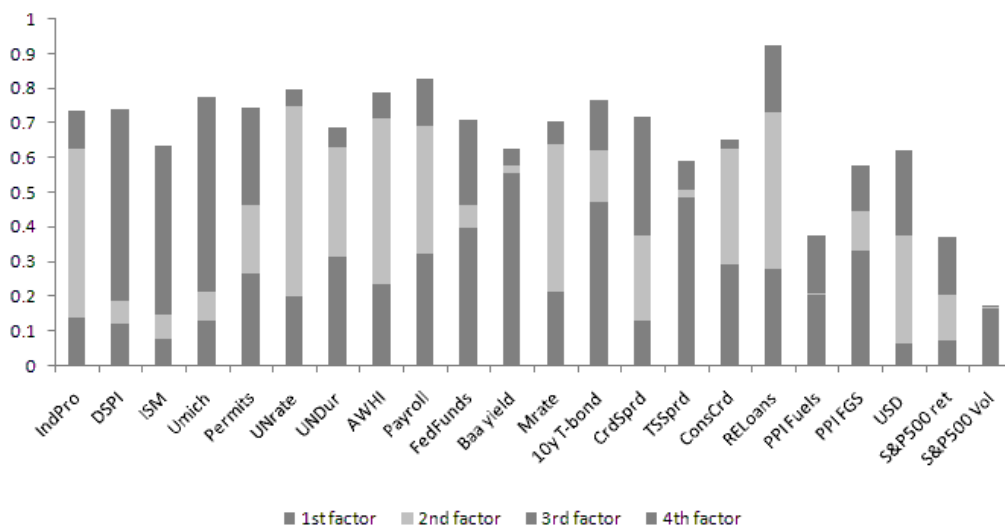
Figure 6: Smoothed Frailty Risk Factor and Industry-group dynamics

The top graph shows the estimated frailty risk factor, which is assumed common to all default counts. The second graph plots six industry-specific risk factors along with asymptotic standard error bands at a 0.05 significance level. High risk factor values imply higher expected default rates.
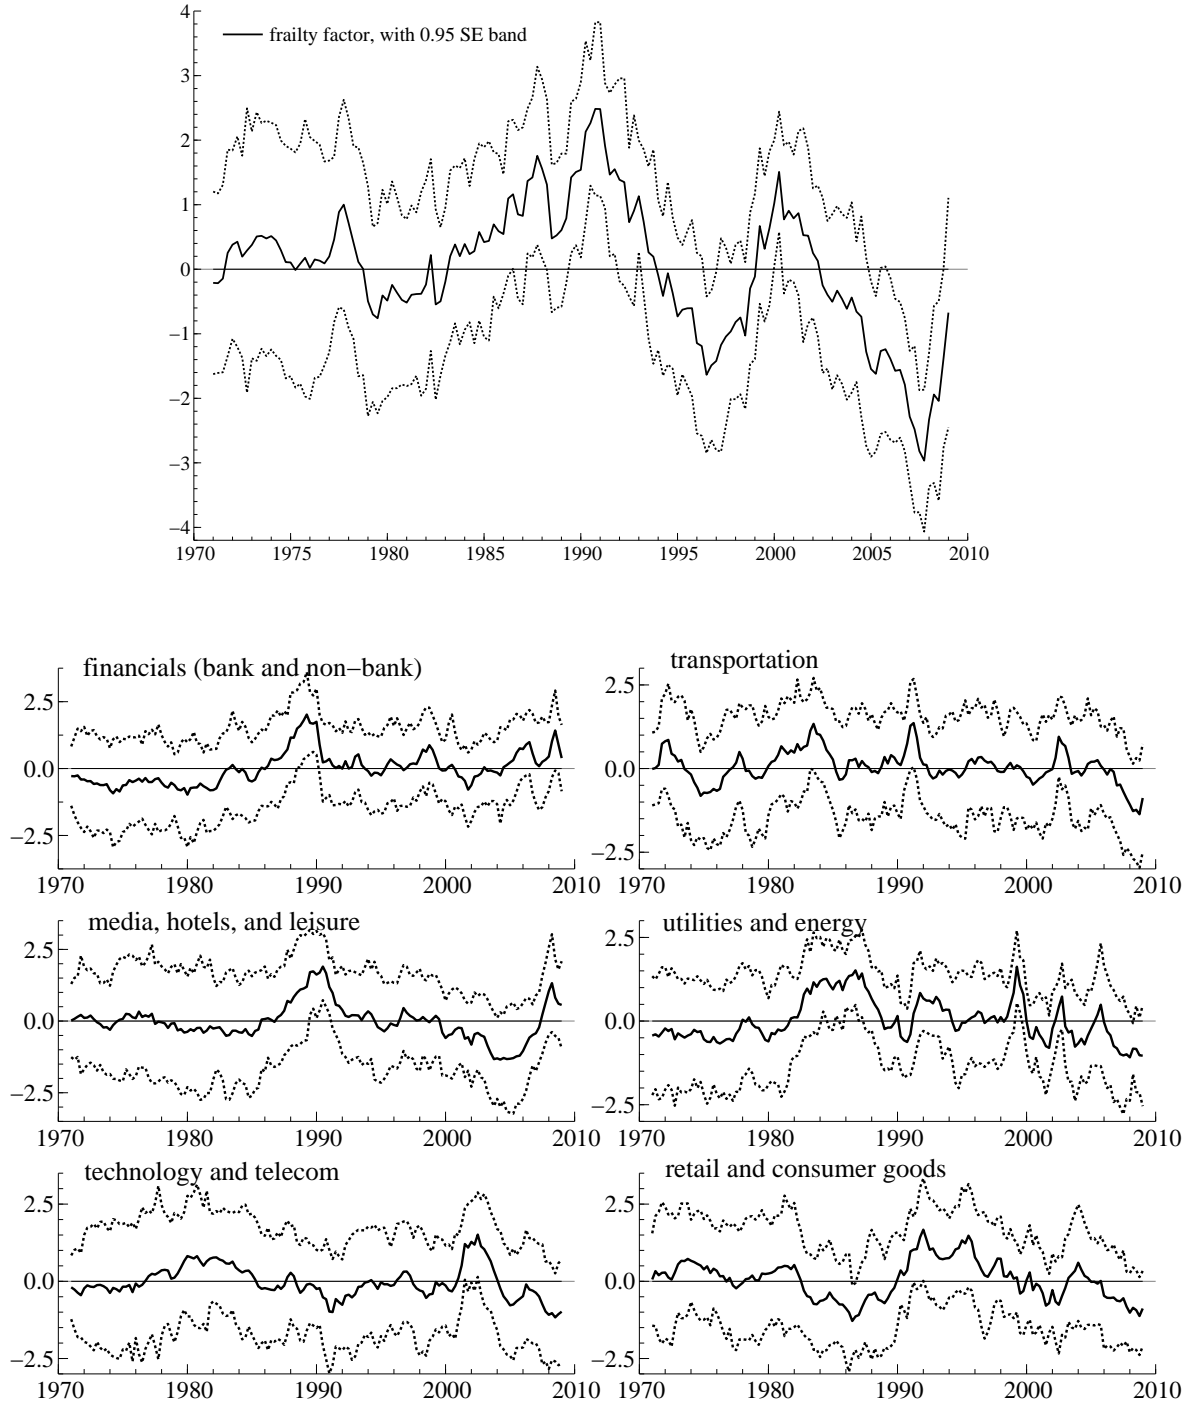
frailty factor implies additional default clustering in these times of stress. On the other hand, the large negative values before the 2007-2009 credit crisis imply defaults that are systematically 'too low' compared to what is implied by macroeconomic and financial data. Both Das et al. (2007) and Duffie, Eckner, Horel, and Saita (2009) ask what effects are captured by the frailty factor. Our estimate in Figure 6 suggests that the frailty factor captures different omitted effects at different times. The frailty factor may capture the positive effects from a high level of asset securitization activity during 2005-2007. In 2001 and 2002, it may capture the negative effects due to the disappearance of trust in accounting information, in response to the Enron and Worldcom scandals. These effects are likely to be important for defaults, but difficult to measure. 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.

Industry factors f_t^i capture deviations of industry-specific dynamics from common variation. For example, we observe industry-specific default stress for financial firms during the U.S. savings and loan crises from 1986-1990, and during the current crisis of 2007-09. Similarly, we observe considerably higher default stress for the technology sector following the 2000/01 asset bubble burst, or for the transportation industry following the 9/11 attack. Lando and Nielsen (2008) observe that it is very difficult to observe evidence for direct default contagion in the Moody's database, based on reading many individual default histories. We confirm this finding to some extent. Our industry factors look more like the industry-specific propagation of economy-wide shocks. For example, the 9/11 shock to the airline industry is visible as a brief spike in the transportation sector at that time. It is difficult to interpret these industry dynamics as contagion. Airlines do not in general lend money to each other, and would gladly take over the remaining market share of a bankrupt competitor. Similarly, the default stress for technology firms in 2001 is clearly visible in the estimated industry-specific risk factor, but is most likely not due to contagion through business links, or indirect contagion through firm's balance sheet data.

Figure 7 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 can 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

Figure 7: 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.
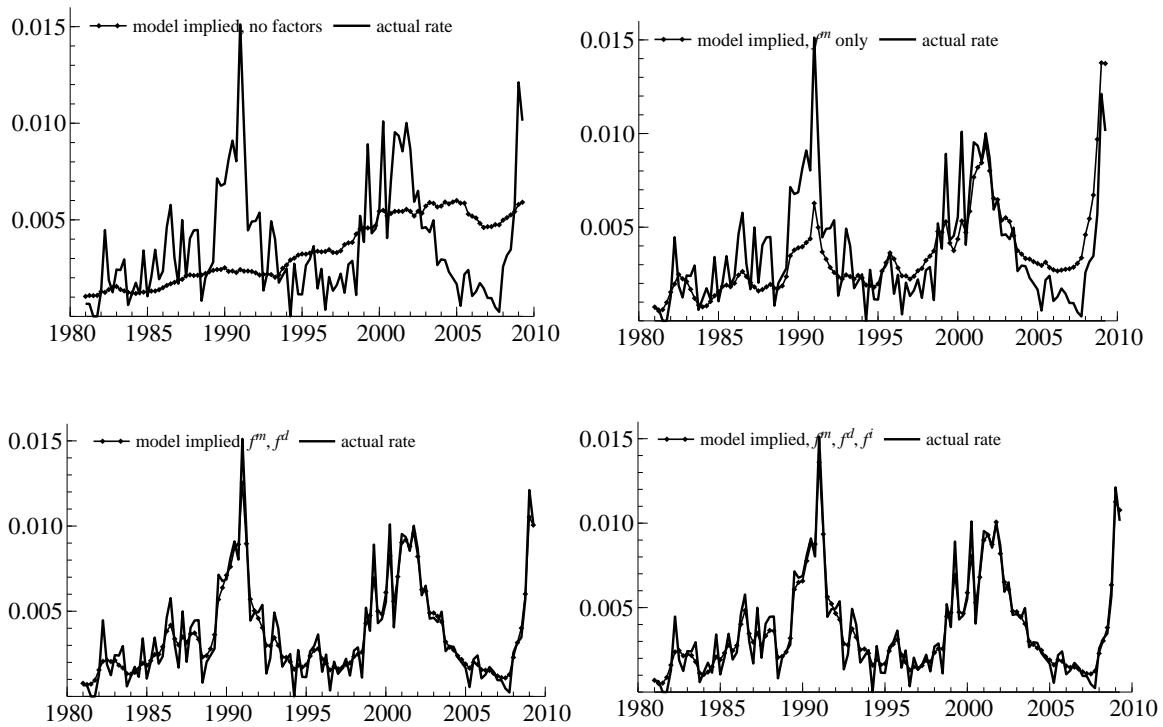
rate variation. The latent frailty dynamics given by f_t^d, however, are clearly required for a good model fit. This holds both in low default periods such as 2002-2007, as well as in high default periods such as 1991. The bottom graphs of Figure 7 indicate that industry-specific developments cancel 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 2.2 to assess which share of default rate volatility is captured by an increasing set of systematic risk factors. The earlier literature on default modeling in the presence of explanatory variables has not addressed this issue in detail.

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. The idiosyncratic risk can be eliminated in a large credit portfolio through diversification. The remaining share of risk, approximately 34%, 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, which may be partly due to default contagion.

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

Table 3: A decomposition of total default risk

The table decomposes total, i.e. 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 data from 1971Q1 to 2009Q1.

Data	Business cycle f_t^c	Frailty risk f_t^d	Industry-level f_t^i	Idiosyncratic $distr.$
Pooled	11.4%	13.9%	8.6%	66.1%
	(33.6%)	(40.9%)	(25.4%)	
Rating groups:				
Aaa-Baa	10.4%	1.1%	6.4%	82.1%
	(58.0%)	(6.3%)	(35.7%)	
Ba	7.1%	7.5%	6.2%	79.2%
	(34.0%)	(36.0%)	(30.0%)	
B	12.5%	22.3%	7.0%	58.2%
	(30.0%)	(53.2%)	(16.8%)	
Caa-C	12.3%	8.9%	12.3%	66.5%
	(36.7%)	(26.5%)	(36.8%)	
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%
High Tech	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 8: Time variation in risk shares

We plot 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.
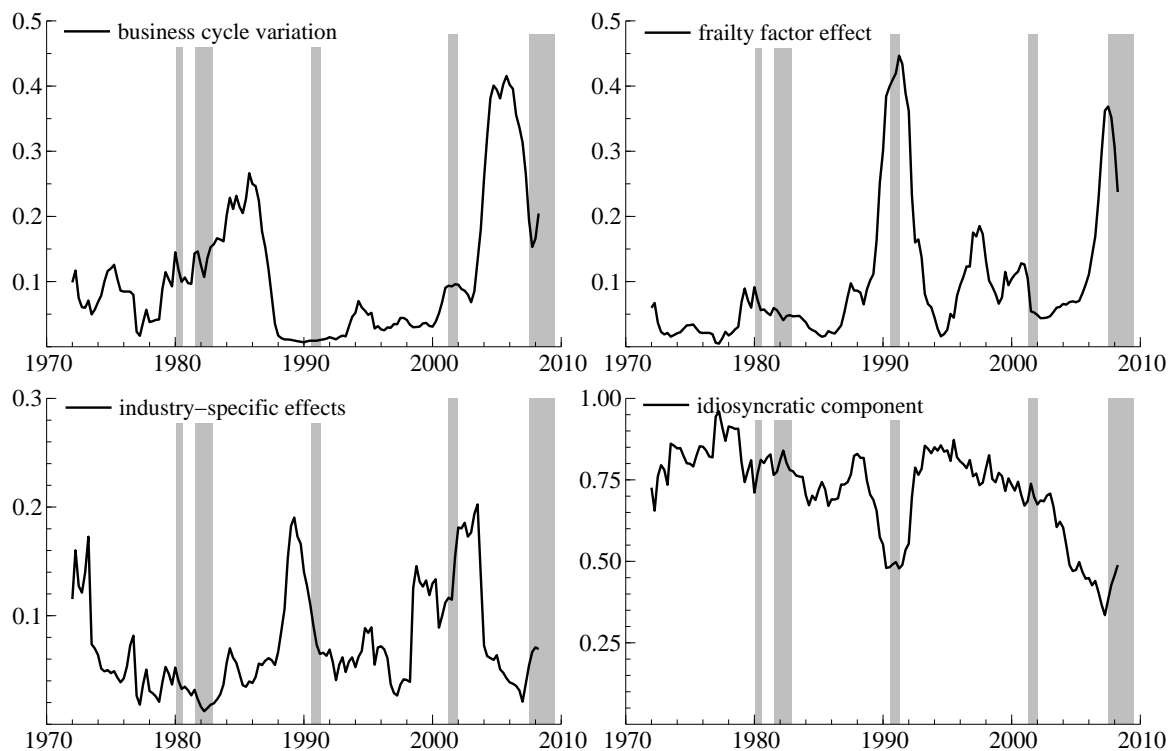
finding can be traced back to two sources. First, the frailty factor loads more heavily on speculative grade firms than investment grade firms. Second, some macro risk factors load on low rating groups also, see Table 2.

Figure 8 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 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

22

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 8 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 rapid correction of this phenomenon over the financial crisis is striking. The eight-quarter rolling R^2 for the macro factors decreases by a factor 2 from 40% to 20% over 2007Q1-2009Q1. Given the rolling window approach, the instantaneous effect may be even higher. The effect is offset by an increased share of explanation due to industry effects (from 2% to 6%) and idiosyncratic risk (from roughly 40% to 50%). Both of these are diversifiable to a lesser or greater extent. The share of explanation due to the frailty factor remains high over the entire crisis period and only decreases towards the end of our sample. 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 the next section.

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.

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.

In case of a default only a certain percentage of the principal is recovered. Rather than using an average recovery rate of around 60%, we assume a stressed recovery rate of 20%. This substantially lower recovery rate accounts for the possible empirical correlation between the probability of default and the recovery rate, see for example Altman, Brady, Resti, and Sironi (2003). Since we are interested in the tail of the loss distribution, the clustering of defaults during periods of low recovery rates is important.

The financial institution uses the reduced form model of Section 3 to determine the appropriate capital buffers. Typically, it picks a high percentile of the predictive loss distribution. Simulating these percentiles is straightforward. First, one uses the filtering methods introduced in Appendix A2 to simulate the current position of the latent systematic risk factors. Next, one 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 (5).

Our 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.

The top panel in Figure 9 contains the credit portfolio loss distribution implied by actual historical default data. This distribution can be compared with the (unconditional) loss distribution implied by three different specifications of our econometric model of 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 bimodality may arise due to the discontinuity in recovered principals in case of default. These qualitative features are confirmed in the top panel of Figure 9.

By comparing the unconditional loss distributions in the top panel of Figure 9, 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

Figure 9: Real vs. model-implied credit portfolio loss distribution

All distribution plots refer to a credit portfolio with uniform loan exposures to Moody's rated firms. The top panel graphs the (unconditional) portfolio 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 remaining plots in the top panel are 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. The bottom panel plots 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. Here, the horizontal axis measures annual losses as fractions of portfolio value.

(b) loss distribution, model with f^m only
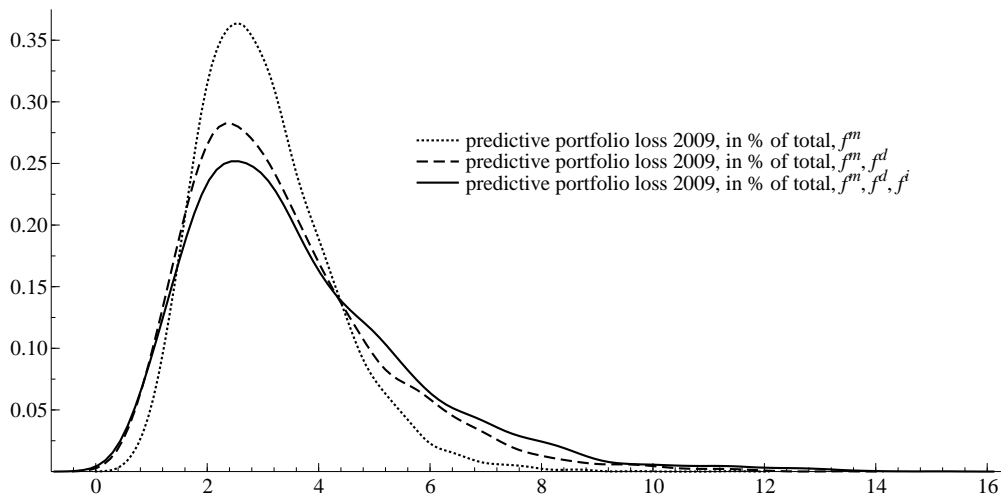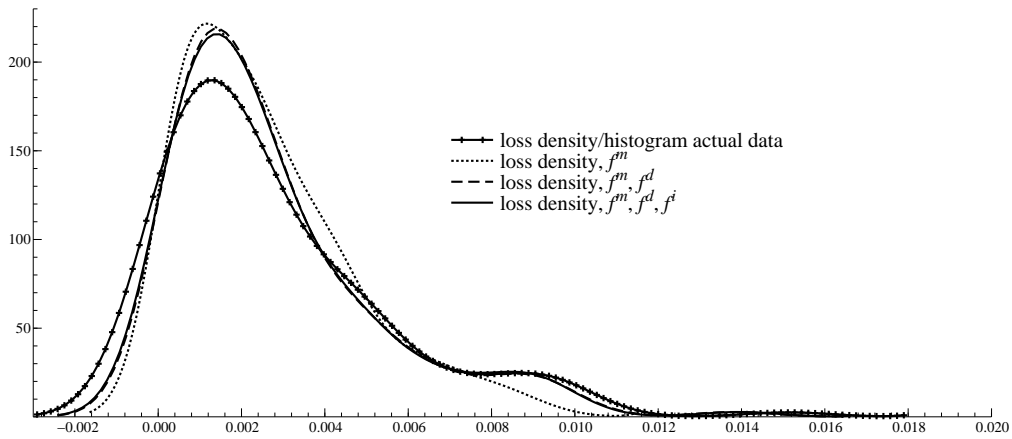
right tail. The loss distributions implied by these models are closer to the actual distribution. The full model is able to reproduce the distributional characteristics of default rates, such as the positive skewness, excess kurtosis, and an irregular shape in the upper tail.

The bottom panel of Figure 9 graphs the simulated predictive credit portfolio loss densities for the years 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. Predicted annual mean losses are roughly comparable at 2.96% and 2.71%, respectively.

4.2   Industry specific risk dynamics

Section 3.2 shows that industry-specific variation accounts for about 17–37% of default rate volatility at the rating and industry level. Industry-specific factors capture the differential impact of each crisis on a given sector. For example, default stress for the banking industry has been high before and during the 1991 and 2008 recessions, but negligible during the 2001 recession. Similarly, while the 2007-2009 crisis is particularly stressful for firms from the financial, manufacturing, and media, hotels, and leisure sector, it is relatively benign on the technology, energy, and transportation sectors.

A specific case illustrates how macro, frailty, and industry-specific dynamics combine to capture industry-level variation in default rates. Figure 10 presents the observed quarterly default fractions rate for the financial sector subsample of the entire Moody's data base. The rates are computed as the percentage of financial sector defaults over the number of firms rated in the financial industry. We compare the observed fractions to the corresponding model-implied rates. 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.

Common variation of defaults based on macroeconomic and financial market covariates captures substantial and overall time-variation in financial sector default rates, see Figure 10. Also, we learn that the frailty factor is of key importance. It captures the overall excessive default activity that is higher before and during the 1991 and 2001 recessions, and substantially lower in the years 2005-2007. The industry-specific factors adjust these common default dynamics to the developments at the sectoral level. The industry-specific factor for financials, as plotted in the second panel of Figure 6, captures the additional sector-specific stress during the banking crisis periods of 1986-1990 and 2007-09. It also

Figure 10: Quarterly time-varying default intensities for financial firms

We plot smoothed estimates of quarterly time-varying default rate 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 graphed against the observed default fractions for financial firms.
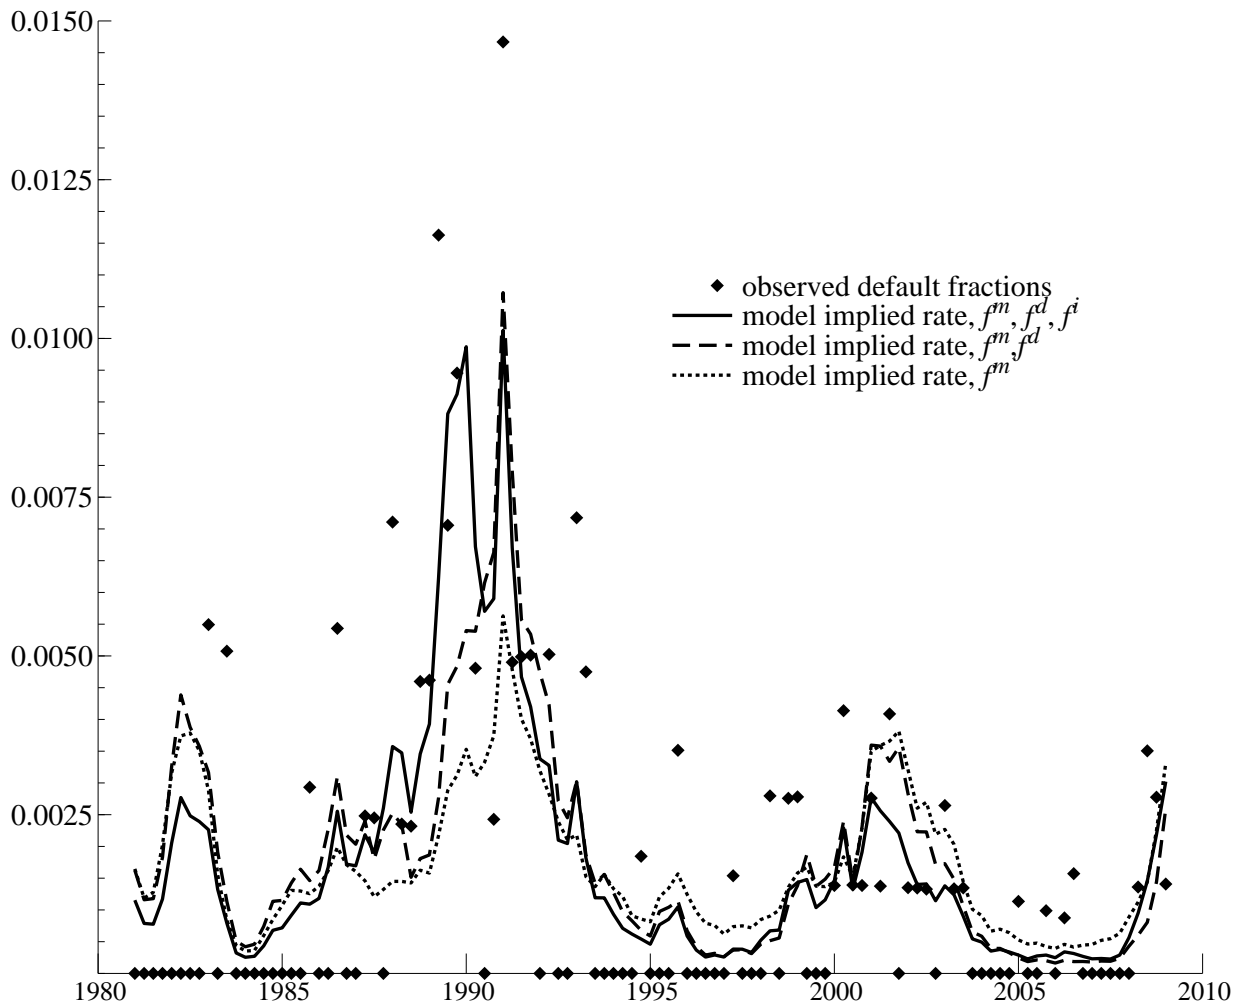
adjusts the default rate (downwards) to the observed lower rates during the 2001 recession.

We conclude that industry factors are important to capture default rates at the industry level. The bottom graphs of Figure 9 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 form a dominant cause for default clustering.

5 Conclusion

We have presented a decomposition of systematic default risk based on a new modeling framework. Observed default counts are modeled jointly with macroeconomic and financial indicators. The resulting panel of continuous and discrete variables is analyzed to investigate the drivers of 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 its three different origins as proposed in the literature. The part due to dependence on common macroeconomic 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 factors in roughly equal proportions. These findings suggest that credit risk management at the portfolio level should account for all three sources of risk simultaneously. In particular, typical industry models that account for macroeconomic dependence only do not account for substantial parts of systematic risk. In particular this holds for commonly adopted stress-testing frameworks, including the 2009 stress-testing exercise for 19 U.S. banks under the Supervisory Capital Assessment Program (SCAP), and the 2010 stress testing exercise for European banks. Omitting frailty and industry factors could be detrimental to a correct assessment of financial system stability.

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.

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.

Appendix A1: exponential family model

We assume throughout that each density $p_j(\cdot)$ is a member of the exponential family. The derivations below can be extended for models outside this class. For the current paper, the exponential family suffices as it contains the normal and binomial densities. We have

$$y_{jt} \sim p_j(y_{jt}|\theta_{jt};\psi), \qquad p_j(y_{jt}|\theta_{jt};\psi) = \exp\left[y_{jt}\theta_{jt} - b_{jt}(\theta_{jt}) + c_{jt}(y_{jt})\right], \qquad \text{(A.10)}$$

and we assume that y_{jt} given θ_{jt} is mutually and serially independent for $j = 1, \ldots, J$ and $t = 1, \ldots, T$. In other words, the dependence in the data set for our MiMe DFM is specified only via the signal θ_{jt} as given by (4) and (2). It follows that $p_j(y_{jt}|\theta_{jt};\psi) = p_j(y_{jt}|\mathcal{F}_t;\psi)$ in (3).

The normal density $p_j(y_{jt}|\theta_{jt};\psi) = \mathrm{N}(\mu_{jt}^*, \sigma_{jt}^2)$ is obtained by having

$$\theta_{jt} = \frac{\mu_{jt}^*}{\sigma^2}, \quad 2b_{jt}(\theta_{jt}) = \theta_{jt}^2 \sigma_{jt}^2 + \ln(2\pi\sigma_{jt}^2), \quad c_{jt}(y_{jt}) = \frac{-y_{jt}^2}{2\sigma_{jt}^2}. \qquad \text{(A.11)}$$

The binomial density is obtained by having

$$\theta_{jt} = \ln\left(\frac{\pi_{jt}}{1 - \pi_{jt}}\right), \quad b_{jt}(\theta_{jt}) = \ln\left(1 + \exp(\theta_{jt})\right), \quad c_{jt}(y_{jt}) = \ln\left(\frac{k_{jt}!}{y_{jt}!(k_{jt} - y_{jt})!}\right). \qquad \text{(A.12)}$$

Appendix A2: 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 (7) 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 $y = (y_1', \ldots, y_T')'$ can be expressed by the joint density of y and $f = (f_1', \ldots, f_T')'$ where f is integrated out, that is

$$p(y;\psi) = \int p(y, f;\psi)\mathrm{d}f = \int p(y|f;\psi)p(f;\psi)\mathrm{d}f, \qquad \text{(A.13)}$$

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

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

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 $k = 1,\ldots,K$. 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

$$p(y;\psi) = \int \frac{p(y,f;\psi)}{g(f|y;\psi)} g(f|y;\psi) \mathrm{d}f = g(y;\psi) \int \frac{p(y|f;\psi)}{g(y|f;\psi)} g(f|y;\psi) \mathrm{d}f. \tag{A.14}$$

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

$$\tilde{p}(y;\psi) = g(y;\psi)M^{-1}\sum_{k=1}^{M} \frac{p(y|f^{(k)};\psi)}{g(y|f^{(k)};\psi)}, \qquad f^{(k)} \sim g(f|y;\psi), \tag{A.15}$$

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

For a practical implementation, the importance density $g(f|y;\psi)$ can be based on the linear Gaussian approximating model

$$y_{jt} = \mu_{jt} + \theta_{jt} + \varepsilon_{jt}, \qquad \varepsilon_{jt} \sim \mathrm{N}(0,\sigma_{jt}^2), \tag{A.16}$$

where mean correction μ_{jt} and variance σ_{jt}^2 are determined in such a way that $g(f|y;\psi)$ is sufficiently close to $p(f|y;\psi)$. It is argued by Shephard and Pitt (1997) and Durbin and Koopman (1997) that μ_{jt} and σ_{jt} can be uniquely chosen such that the modes of $p(f|y;\psi)$ and $g(f|y;\psi)$ with respect to f are equal, for a given value of ψ.

To simulate values from the importance density $g(f|y;\psi)$, the simulation smoothing method of Durbin and Koopman (2002) can be applied to the approximating model (A.16). For a set of M draws of $g(f|y;\psi)$, the evaluation of (A.15) relies on the computation of $p(y|f;\psi)$, $g(y|f;\psi)$ and $g(y;\psi)$. Density $p(y|f;\psi)$ is based on (3), density $g(y|f;\psi)$ is based on the Gaussian density for $y_{jt} - \mu_{jt} - \theta_{jt} \sim \mathrm{N}(0,\sigma_{jt}^2)$ (A.16) and $g(y;\psi)$ can be computed by the Kalman filter applied to (A.16), see Durbin and Koopman (2001).

The likelihood function can be evaluated for any value of ψ. For a given set of random numbers from which factors are simulated from $g(f|y;\psi)$, we maximize the likelihood (A.15) with respect to ψ.

Furthermore, we can estimate the latent factors f_t via importance sampling. It can be shown that

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

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

$$\tilde{f} = \sum_{k=1}^{M} w_k \cdot f^{(k)} \left/ \sum_{k=1}^{M} w_k, \right.$$

with $w_k = p(y|f^{(k)};\psi)/g(y|f^{(k)};\psi)$, and $f^{(k)} \sim g(f|y;\psi)$. Similarly, the standard errors s_t of $\tilde{f}_t$ can be estimated by

$$s_t^2 = \left(\sum_{k=1}^{M} w_k \cdot (f_t^{(k)})^2 \left/ \sum_{k=1}^{M} w_k \right. \right) - \tilde{f}_t^2,$$

with $\tilde{f}_t$ the tth elements of $\tilde{f}$. Conditional mode estimates of the factors are given by

$$\bar{f} = \operatorname{argmax} p(f|y; \psi), \tag{A.17}$$

and indicate the most probable value of the factors given the observations. They are obtained as a by-product when matching the modes of densities $p(f|y; \psi)$ and $g(f|y; \psi)$.

References

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.

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.

Das, S., D. Duffie, N. Kapadia, and L. Saita (2007). Common Failings: How Corporate Defaults Are Correlated. *The 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. *Biometrica 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.

Lando, D. and M. S. Nielsen (2008). 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 bankrupcy announcements. *Journal of Financial Economics 32*, 45–60.

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. *Biometrica 84*, 653–67.