

Credit Risk and State Space Methods

ISBN 978 90 3610 212 4

Cover design: Crasborn Graphic Designers bno, Valkenburg a.d. Geul

This book is no. **488** of the Tinbergen Institute Research Series, established through cooperation between Thela Thesis and the Tinbergen Institute. A list of books which already appeared in the series can be found in the back.

VRIJE UNIVERSITEIT

Credit Risk and State Space Methods

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad Doctor aan
de Vrije Universiteit Amsterdam,
op gezag van de rector magnificus
prof.dr. L.M. Bouter,
in het openbaar te verdedigen
ten overstaan van de promotiecommissie
van de faculteit der Economische Wetenschappen en Bedrijfskunde
op maandag 17 januari 2011 om 13.45 uur
in de aula van de universiteit,
De Boelelaan 1105

door

Bernd Schwaab

geboren te Frankfurt am Main, Duitsland.

promotoren: Andre Lucas

Siem Jan Koopman

Acknowledgements

Nobody makes it alone in work and life. This is particularly true for a PhD candidate. First and foremost, I thank my supervisors André Lucas and Siem Jan Koopman. Thank you for taking me on as a PhD student in February 2007. That was a risk, since at that point I had heard about co-integration and economics, but knew little about credit risk and state space methods. I have had my share of success because I have been standing on your shoulders since that time. You create lots of value for other people. No wonder you are busy. Thanks for being accessible throughout these years.

As of June 2010, the contents of this thesis have been presented on 28 occasions, in 15 different cities, in 6 different countries. The presentations included prestigious meetings in Amsterdam, Atlanta, Chicago (twice), Milan, San Francisco, and Venice, among others, and more are coming up. Many researchers in credit risk and econometrics at those meetings had a share in pointing me towards questions for which people have intuition, but nobody knows exactly. I thank VU University Amsterdam and Tinbergen Institute for providing funding for presentations at those meetings.

I am grateful to Prof. Drew Creal, now at the University of Chicago Booth. Thanks for sharing your insight and econometric intuition already in Amsterdam, and for later having me as a visiting PhD student at Booth. I remember my ‘own’ office on the faculty floor, a favor usually extended to only the most distinguished visitors. I recall meeting many people way above my pay grade. While only a few weeks, this visit had a definite impact. I also thank the C. Willems Stichting for providing funding for that visit.

Many good friendships were forged during my first year at Tinbergen Institute (TI), when the heat was on. At TI, I thank fellow students and staff Jaap Abbring, Jonneke Bolhaar, Dion Bongaerts, Andrei Dubovik, David Hollanders, Sumedha Gupta, Kobus de Hoop, Arianne de Jong, Marloes Lammers, Pan Lei, Yang Nan, Petr Sedlacek, Marcelo Tyszler, Melinda Vigh, Nick Vikander, Marcel Vorage, Michel van der Wel, Ronald Wolthoff, among many others, for making Tinbergen Institute a great place to work and hang out. At VU University I thank my office mates Liv, Mahmoud, Ting, and Sunny, and the econometricians Bahar, Borus, Brian, Irma, and Taoying, for their

companionship and chats over lunch during the last three years.

At Tinbergen I was able to build on my prior education at Clark University in Worcester (MA). I thank the Fulbright Commission in Berlin for funding during my first year in the U.S., and the Economics Department at Clark for funding the second year. I am grateful to my then academic advisor and mentor Daniel Bernhofen for guidance and for helping facilitate the transition to Tinbergen Institute.

Sometimes the best research ideas come while enjoying other activities. I thank my Salsa dance partners Marloes, Debby, and Juul, for great times in Cantinero and other places. I also thank pastors Steve and Lizby Warren, James and Corrie Herbertson, and Anne and Jedidja Borkent from C3 church in Amstelveen for imparting some of their wisdom. There is power in having your non-dull alive church to go to. Finally, I thank Lubna for her amazing ability to take my mind off things. Last, but certainly not least, I thank my mom and dad and extended family, for their support and for just being there.

Bernd Schwaab

Amsterdam

June 2010

Contents

Acknowledgements	v
Contents	vii
1 Introduction	1
1.1 Motivation	1
1.2 Why state space methods for credit risk?	3
1.3 State space models in a nutshell	4
1.4 Outline of the thesis	6
2 Modeling Frailty-correlated Defaults Using Many Macroeconomic Co- variates	9
2.1 Introduction	9
2.2 The econometric framework	12
2.3 The financial framework	14
2.4 Estimation using state space methods	16
2.4.1 Estimation of the macro factors	16
2.4.2 The factor model in state space form	17
2.4.3 Parameter estimation and signal extraction	18
2.4.4 Simulation experiments	21
2.5 Estimation results and forecasting accuracy	22
2.5.1 Data	22
2.5.2 Macro and contagion factors	27
2.5.3 Model specification	30
2.5.4 Empirical findings	31
2.5.5 Interpretation of the frailty factor	33
2.5.6 Out of sample forecasting accuracy	36
2.6 Conclusion	40

3	Mixed measurement dynamic factor models	41
3.1	Introduction	41
3.2	Mixed-measurement dynamic factor model	43
3.2.1	Model specification	43
3.2.2	Estimation via Monte Carlo maximum likelihood	44
3.2.3	Estimation of the factors	47
3.2.4	Collapsing observations	48
3.2.5	Missing values due to mixed frequencies and forecasting	49
3.3	Mixed measurement generalized autoregressive score models	50
3.3.1	Model specification MM-GAS	51
3.3.2	Maximum likelihood estimation	51
3.4	Bayesian inference	52
3.4.1	Sampling the latent factors	53
3.4.2	Sampling factor loadings and autoregressive parameters	53
3.5	Intertwined credit and recovery risk	54
3.5.1	Data and mixed measurement model equations	54
3.5.2	Major empirical findings	56
3.5.3	Out of sample evaluation	60
3.6	Conclusion	65
4	Macro, frailty, and contagion effects in defaults: lessons from the 2008 credit crisis	69
4.1	Introduction	69
4.2	A joint model for default, macro, and industry risk	72
4.2.1	The mixed measurement dynamic factor model	73
4.2.2	Decomposition of non-Gaussian variation	75
4.3	Empirical findings for U.S. default and macro data	77
4.3.1	Major empirical results	82
4.3.2	Total default risk: a decomposition	89
4.4	Implications for risk management	92
4.4.1	The frailty factor	93
4.4.2	Industry specific risk dynamics	95
4.5	Conclusion	97
5	A diagnostic framework for financial systemic risk assessment	101
5.1	Introduction	101
5.2	The econometric framework	104

5.2.1	Conditioning variables for capital requirements	106
5.2.2	Systematic default risk indicator	108
5.2.3	Magnitude of frailty effects	109
5.3	Data	110
5.4	Major empirical results	115
5.4.1	Macro, frailty, and contagion effects	116
5.4.2	Common default stress	121
5.5	Early warning signals	125
5.6	Conclusion	129
6	Conclusion	131
	Bibliography	135
	Samenvatting (Summary in Dutch)	143

Chapter 1

Introduction

1.1 Motivation

The writing of this thesis coincided with the financial crisis of 2007–09. The causes of the crisis were complex and varied, and have by now been analyzed and documented, see e.g. the de Larosiere report (2009) for the European Union, the Geneva Report on the World Economy of Brunnermeier, Crocket, Goodhart, Persaud, and Shin (2009), and the G30 Report of the Group of Thirty, see Volcker et al. (2009). These reports reach similar broad conclusions about the causes of, and main lessons from, the financial crisis. They uncover important shortcomings in current financial regulation, risk management practises at the firm level, and more generally our limited understanding of risk dynamics and the interplay of credit and macroeconomic conditions.

Among many other issues, regulatory frameworks around the world appear to pay too little attention to systemic risk. An important subset of financial systemic risk is systematic credit risk. Systematic credit risk is the part of portfolio risk that does not average out in the cross section due to diversification effects. Such dependence undermines positive effects from diversification at the firm level, and may further lead to a ‘fallacy of composition’ at the systemic level, as described e.g. in Brunnermeier et al. (2009). Essentially, traditional risk-based capital regulation alone may underestimate systemic risk by neglecting the macro impact of banks reacting in unison to a shock. As a result, measurement of financial systemic risk must necessarily include an assessment of systematic credit risk conditions and its underlying sources. Changes in systematic risk factors can account for pronounced default rate volatility at the firm level, and explain observed default clustering at the aggregate level. Such default clustering is one of the main risks in the banking book, and visible both at the aggregate as well as the disaggregated (industry and rating group) level.

Credit defaults may be assumed to be dependent in the cross section for at least two reasons. First, default dependence arises from exposure to common observed and unobserved risk factors. For example, all firms in an economy are subject to the same business cycle conditions, monetary policy, fiscal policy, financial market prices, swings of optimism and pessimism, trust in the accuracy of accounting numbers, access to credit, etc. These factors combine to correlate firms' defaults at a given time, and may or may not be easily observed. Second, business and other contractual links may give rise to default chains through default contagion. Contagion refers to the phenomenon that a defaulting firm may weaken other firms with which it has direct business and contractual links. Contagion may therefore generate additional default dependence in the cross section, in particular at the industry level. In the main chapters to follow, we control for either source of default clustering by using observed and/or unobserved risk factors.

Current levels of systematic credit risk, while important for defaults and financial systemic risk assessment, are latent since they are simply not observed. Dynamic latent processes are known to econometricians as 'unobserved components'. Given the time-varying nature of latent systematic risk, it is natural to use unobserved component techniques based on state space methods for estimating the location and dynamics of risk factors from observed data. The non-Gaussian nature of credit risk data, usually involving either (discrete) event counts or (nonnegative, continuous) transition spells, implies that nonlinear non-Gaussian models are required.

So then, why do corporate defaults cluster over time? Which sources contribute to default clustering, and to which extent? Is variation in observed macroeconomic and financial data sufficient to explain default rate volatility, or is there evidence for an additional latent (frailty) component driving defaults? If so, what does the frailty factor capture? To what extent are contagion dynamics important at the portfolio level? From an econometric perspective, which framework permits the estimation of parameters and latent factors from large dimensional panels in reasonable amounts of time, such as less than one hour? How can we model default conditions jointly with other data of interest, such as loss-given-default, when observations come from different families of parametric distributions? Are there less complex alternatives to models in state space form? From a policy perspective, how does systematic credit risk relate to overall financial systemic risk? Can policy makers obtain default cycle measurements, and warning signals for financial stability? Are frailty dynamics important for countries other than the U.S.? Are there benefits from tracking default conditions around the globe? The following chapters will address these and other related questions.

1.2 Why state space methods for credit risk?

Why use latent dynamic factor modeling techniques based on state space methods to address problems from credit risk? There are at least three reasons.

First, and despite substantial amounts of research on the topic, credit risk practitioners still have an incomplete understanding about which processes drive the common (or systematic) variation in corporate default hazard rates. To capture such shared dynamics empirically, it is unclear whether one should, for example, include as right hand side variables the growth rates of real gross domestic product, a short term interest rate, the trailing one-year return on a broad equity index, equity volatility, changes in unemployment rate, the yield spread of corporate bonds over treasuries, the term structure, all of the above and many more, or a small subset of these covariates as observed common risk factors. Different factors are found to be significant in different studies, with surprisingly little overlap. Including a latent risk factor in addition to observed data can be seen as ‘**insurance**’ against dynamic model misspecification due to omitted relevant variables and other missing systematic effects. This is important, since the omission of a systematic component causes a downward bias in the estimation of default rate volatility, and in the model-implied probability of extreme portfolio losses. In turn, dynamic latent components are most easily handled for models in state space form.

Second, unobserved latent (frailty) factors are a convenient device to capture **excess default clustering** in default data. Recent research indicates that observed macroeconomic and financial variables and firm-level information are not sufficient to capture the large degree of default clustering in observed corporate default data. Credit risk researchers often reject the joint hypothesis of (i) well-specified default intensities in terms of observed financial and macroeconomic variables and firm-specific information and (ii) the conditional independence (doubly stochastic default times) assumption. This is bad news for practitioners, since virtually all current credit risk models build on conditional independence. Frailty models allow to retain the conditional independence assumption by capturing default dependence above and beyond what is implied by observed risk factors.

Third, the leftover dynamics in standard models of portfolio credit risk may be **interesting and useful** in their own right. In *The Black Swan*, Taleb (2007) refers to the many unread books in the writer Umberto Eco’s library as a metaphor for the notion that an appreciation of the unknown unknown is important, and that the most useful information is often outside the realm of regular expectation. Models with frailty effects essentially allow us to put structure on what is *missing*. Estimated frailty effects are informative about how and when standard models tend to go wrong. Chapter 5 suggests

that the magnitude of estimated frailty effects, and thus the extent to which credit and business cycle conditions may decouple from each other, could have served as a warning signal for a macro-prudential policy maker in charge of financial stability.

The apparent difficulties in attributing common variation in default hazard rates (or rating transition intensities) to observed risk factors are to be expected if default and macroeconomic conditions are related but inherently different processes. This is a recurring finding in this thesis. For example, in Chapter 2 we find a large and significant role for a dynamic residual (frailty) component even after controlling for more than eighty percent of the variation in more than hundred macroeconomic and financial covariates, as well as industry level contagion dynamics and equity information. Chapter 4 finds that observed macro and financial market factors account for only 30–60% of systematic default risk. Chapter 5 presents evidence that latent residual dynamics are important also for non-U.S. data. Latent dynamic factor techniques are required for such assessments, since, obviously, neither the business cycle nor levels of systematic credit risk are directly observed.

If default conditions can significantly and persistently diverge from what is implied by macroeconomic and financial market covariates, then inference on the default cycle using business cycle measurements is at best suboptimal and at worst systematically misleading. Using a few observed macros to predict corporate defaults out-of-sample is then also little more than an exercise in wishful thinking and possibly self-deception about the incurred level of risk. If financial industry models based on similar standard risk factors only are used to calculate capital buffers, they will tend to be wrong all at the same time. By contrast, dynamic latent factor models based on state space methods allow us to improve the out of sample forecasting accuracy for corporate default hazard rates (Chapter 2), model systematic effects across different types of observations from different families of parametric distributions (Chapter 3), assess in more detail which sources drive default clustering and to which extent (Chapter 4), and obtain a diagnostic framework and warning signals for financial systemic stability (Chapter 5). A brief introduction to state space methods is provided next.

1.3 State space models in a nutshell

This section provides a (very) brief and non-technical discussion of models in state space form. Chapters 2 to 5 provide additional detail when state space methods are applied to specific models. The discussion is based on Durbin and Koopman (2001) and Jungbacker and Koopman (2007).

The state space form provides a unified representation of a wide range of linear Gaussian and nonlinear non-Gaussian time series models. Examples of linear Gaussian models in state space form are autoregressive moving average (ARMA) models, time-varying regression models, dynamic linear models, and structural unobserved components time series models, see e.g. Harvey (1993) and West and Harrison (1997). Examples of models that are not both linear and Gaussian include stochastic volatility models, more complicated unobserved component models, and models for defaults and rating transitions, see e.g. Koopman and Lucas (2008), Koopman, Lucas, and Monteiro (2008), and Lee (2010). The state space form consists of a measurement, signal, and state equation, and can be stated as

$$y_t \sim p(y_t|\theta_t) \tag{1.1}$$

$$\theta_t = c_t + Z_t\alpha_t \tag{1.2}$$

$$\alpha_{t+1} = d_t + T_t\alpha_t + R_t\eta_t, \quad \eta_t \sim \text{NID}(0, H_t), \tag{1.3}$$

$t = 1, \dots, T$, where the vector of observations y_t depends on a vector of signals θ_t , the signal θ_t is an affine function of elements in the latent state vector α_t , which in turn evolves over time as a first order Markov process driven by normally distributed and serially uncorrelated innovations η_t . The state vector may contain unobserved stochastic processes such as latent dynamic factors and unknown fixed effects. The initial state α_1 is assumed to be random with a given initial mean and (possibly diffuse) variance matrix. System matrices Z_t , T_t , R_t and H_t and intercepts c_t and d_t may vary over time and may depend on unknown coefficients. Unknown parameters from the system matrices and intercepts are collected in a vector ψ , and are usually estimated by maximum likelihood. The state space form allows for two independent sources of error, given by $p(y_t|\theta_t)$ in the measurement equation, and by η_t in the state equation. Consequently, the data may be observed with noise, and the elements of the state vector may be serially correlated.

If the observations y_t in (1.1) are Gaussian, parameter and latent state vector estimation is relatively straightforward. The observation equation becomes linear

$$y_t = \theta_t + G_t\epsilon_t, \quad \epsilon_t \sim \text{NID}(0, I),$$

where G_t is a matrix which may depend on ψ . The model log-likelihood can be obtained efficiently using one run of the Kalman Filter (KF). The KF recursions efficiently compute minimum mean squared error predictions, prediction errors, and prediction error variances

associated with the state and observation vectors given parameters and past observations. The associated smoother (KFS) uses the output of the KF to compute the conditional expectation of the states given the complete sample y_1, \dots, y_T .

If observations y_t are not Gaussian, but, for example, come from another member of the exponential family of densities, then the measurement equation (1.1) is nonlinear in the signal θ_t . Important densities such as $p(\alpha|y)$ and the model likelihood $p(y; \psi)$ are not available in closed form. This obviously hinders a likelihood-based analysis of the model. Simulation-based techniques are almost always required.

To overcome the problem that the model log-likelihood does not exist in closed form for non-Gaussian models in state space form, this thesis makes extensive use of Monte Carlo maximum likelihood techniques based on importance sampling. The details are given in later chapters, but we essentially follow a similar recipe each time. The log-likelihood is maximized by, at each evaluation of the log-likelihood, integrating out the unobserved components in α_t from their joint density with the observations. In each evaluation, we determine the conditional mode of the state given the observations for the non-Gaussian model. We then find the linear Gaussian model with the same conditional mode given the observations by iterated use of the Kalman filter and smoother. We use the conditional density of the state given artificial observations $g(\alpha|\tilde{y})$ for this model as the importance density. Since this density is Gaussian, we can draw random samples from it in a relatively straightforward and efficient way using a simulation smoother. Given maximum likelihood estimates, conditional mean and variance estimates of the elements in the state vector can similarly be obtained from mode estimates.

1.4 Outline of the thesis

This thesis consists of four main chapters. In each chapter we apply latent dynamic factor modeling techniques, usually in a state space framework, to a credit risk modeling problem at hand. Due to the non-Gaussian nature of credit risk data, we rely on non-Gaussian models in state space form. Each chapter is self-contained and can be read independently. Therefore, each chapter ends with a discussion of its contribution.

Chapter 2 ("*Modeling frailty-correlated defaults using many macroeconomic covariates*") is based on Koopman, Lucas, and Schwaab (2008). In this chapter, we propose a new econometric framework for estimating and forecasting the default intensities of corporate credit subject to observed and unobserved risk factors. The model combines common factors from macroeconomic and financial covariates with an unobserved latent (frailty) component for discrete default counts, observed contagion factors at the industry

level, and standard risk measures such as ratings, equity returns, and volatilities. In an empirical application, we find a large and significant role for a dynamic frailty component even after controlling for more than eighty percent of the variation in more than hundred macroeconomic and financial covariates, as well as industry level contagion dynamics and equity information. We emphasize the need for a latent component to prevent the downward bias in estimated default rate volatility at the rating and industry levels and in estimated probabilities of extreme default losses on portfolios of U.S. debt. The latent factor does not substitute for a single omitted macroeconomic variable. We argue that it captures different omitted effects at different times. We also provide empirical evidence that default and business cycle conditions depend on different processes. In an out-of-sample forecasting study for point-in-time default probabilities, we obtain mean absolute error reductions of more than forty percent when compared to models with observed risk factors only. The forecasts are relatively more accurate when default conditions diverge from aggregate macroeconomic conditions.

Chapter 3 ("*Mixed measurement dynamic factor models*") is based on Creal, Schwaab, Koopman, and Lucas (2010). In this chapter, we propose a new latent dynamic factor model framework for mixed-measurement mixed-frequency panel data. Time series observations may come from different families of parametric distributions, may be observed at different frequencies, and exhibit common dynamics and cross sectional dependence due to shared exposure to latent dynamic factors. As the main complication, the likelihood does not exist in closed form for this class of models. We therefore present different approaches to parameter and factor estimation in this framework. First, assuming a factor structure for location parameters yields a parameter driven model that can be cast into state space form. Parameters and factors estimation is then accomplished by Monte-Carlo maximum likelihood based on importance sampling. Second, we propose a less complex observation driven alternative to the parameter driven original model, for which the likelihood exists in closed form. Finally, parameter and factor estimates can be obtained by Markov Chain Monte Carlo. We use the new mixed-measurement framework for the estimation and forecasting of intertwined credit and recovery risk conditions for US Moody's-rated firms from 1982 - 2008. The joint model allows us to construct predictive (conditional) loss densities for portfolios of bank loans and corporate bonds in the presence of non-standard sources of credit risk such as systematic frailty effects and systematic recovery risk.

Chapter 4 ("*Macro, frailty, and contagion effects in defaults: lessons from the 2008 credit crisis*") is based on Koopman, Lucas, and Schwaab (2010). Default clustering is one of the main risks in the banking book. Three explanations have been proposed for such clustering. First, defaults may covary with business cycle and financial market conditions.

Second, defaults can have their own frailty dynamics separate from the business cycle. Third, there may be industry-specific dynamics including contagion that give rise to default clusters. We develop a new integrated empirical modeling framework that allows us to disentangle, quantify, and test these competing explanations. Using US firm data from Moody's over the period 1971–2009, we find that systematic risk factors account for roughly one third of observed default rate volatility. Observed macro and financial market factors, in turn, account for 30–60% of this systematic default risk. Consequently, credit risk models that only control for macro conditions leave out a substantial share of systematic default rate variation. The remainder of systematic default risk is captured by frailty effects, closely followed by industry effects. The frailty components are particularly relevant around times of stress. We find that in the years leading up to the 2008 financial crisis default risk has been systematically too low compared to what one would expect based on a wide range of macro variables. The framework may thus also provide a tool to detect systemic risk build-up in the economy.

Chapter 5 (*"A diagnostic framework for financial systemic risk assessment"*) is based on Schwaab, Lucas, and Koopman (2010). A macro-prudential policy maker can manage financial systemic risks only if such risks can be reliably assessed. To this purpose we propose a large-scale modeling framework for the measurement of international macroeconomic and credit risk conditions. The model can be used as a diagnostic tool to track the evolution and composition of credit risk and default clustering around the globe. We present an indicator to summarize common default stress for a given set of firms. When applied to financial firms, we obtain a straightforward measure of unobserved financial systemic risk. In an empirical analysis of worldwide credit data for more than 12,000 firms in four broad economic regions, we find that default conditions can significantly and persistently decouple from what is implied by macroeconomic and financial data due to latent frailty risk. We then suggest that the magnitude of estimated frailty effects can serve as an early warning signal for macro-prudential policy makers. Frailty effects have been pronounced during bad times, such as the savings and loan crisis in the U.S. leading up to the 1991 recession, and exceptionally good times, such as the years 2005–07 leading up to the recent financial crisis, when defaults were much lower than implied by macro and financial data in many parts of the world.

Chapter 6 (*"Conclusion"*) summarizes the main results and concludes the thesis.

Chapter 2

Modeling Frailty-correlated Defaults Using Many Macroeconomic Covariates

2.1 Introduction

Recent research indicates that observed macroeconomic variables and firm-level information are not sufficient to capture the large degree of default clustering in observed corporate default data. In an important study, Das, Duffie, Kapadia, and Saita (2007) reject the joint hypothesis of (i) well-specified default intensities in terms of observed macroeconomic variables and firm-specific information and (ii) the conditional independence (doubly stochastic default times) assumption. This is bad news for practitioners, since virtually all current credit risk models build on conditional independence.

Excess default clustering is often attributed to frailty and contagion. The frailty effect captures default dependence that cannot be captured by observed macroeconomic and financial data. In the econometric literature the frailty effects are usually modeled by an unobserved risk factor, see McNeil and Wendin (2007), Azizpour and Giesecke (2008), Koopman, Lucas, and Monteiro (2008), Koopman and Lucas (2008), and Duffie, Eckner, Horel, and Saita (2009). When a model for discrete default counts contains dynamic latent components, the likelihood function is not available in closed form and advanced econometric techniques based on simulation methods are required. For this reason McNeil and Wendin (2007) and Duffie et al. (2009) employ Bayesian inference methods, while Koopman et al. (2008) and Koopman and Lucas (2008) rely on a Monte Carlo maximum likelihood approach.

In addition to frailty effects, contagion dynamics offer another source of default clus-

tering. Contagion refers to the phenomenon that a defaulting firm can weaken the firms in its network of business links, see Giesecke (2004) and Lando and Nielsen (2008). Such business links are particularly relevant at the industry level through supply chain relationships, see Lang and Stulz (1992), Jorion and Zhang (2007b), and Boissay and Gropp (2010).

In this paper we develop a practical and feasible econometric framework for the measurement and forecasting of point-in-time default probabilities. The underlying economic model allows for default correlations that originate from macroeconomic and financial conditions, frailty risk and contagion risk. The model is aimed to support credit risk management at financial institutions. It may also have an impact on the assessment of systemic risk conditions at (macro-prudential) supervisory agencies such as the new European Systemic Risk Board (ESRB) for the European Union, and the Financial Services Oversight Council (FSOC) for the United States. Time-varying default risk conditions contribute to overall financial systemic risk, and an assessment of the latter requires estimation of the former.

We present three contributions to the econometric credit risk literature. First, we show how a nonlinear non-Gaussian panel data model for discrete default counts can be combined with an approximate dynamic factor model for continuous macroeconomic time series data. The resulting model inherits the best of both worlds. A linear Gaussian factor model permits the use of information from large arrays of relevant predictor variables for the modeling of defaults. The nonlinear non-Gaussian panel data model in state space form allows for unobserved frailty effects, accommodates the cross-sectional heterogeneity of firms, and handles missing values that arise in count data at a highly disaggregated level. In effect, our model combines a non-Gaussian panel specification with a dynamic factor model for continuously valued time series data as used in, for example, Stock and Watson (2002b). Parameter and factor estimation are achieved by adopting a maximum likelihood framework and using importance sampling techniques derived for multivariate non-Gaussian models in state space form, see Durbin and Koopman (1997, 2001) and Koopman and Lucas (2008). The resulting framework allows us to estimate a large dimensional econometric model for time-varying default conditions, which accommodates 112 time series of disaggregated default counts and more than 100 macroeconomic and financial covariates, in only 20 - 90 minutes on a standard desktop PC. The computational speed and model tractability allows us to conduct repeated out-of-sample forecasting experiments, where parameters and factors are re-estimated based on expanding sets of data.

Second, in an empirical study of U.S. default data from 1981Q1 to 2009Q4, we find a

large and significant role for a dynamic frailty component even after taking into account more than 80% of the variation from more than 100 macroeconomic and financial covariates, while controlling for contagion at the industry level as well as standard measures of risk such as ratings, equity returns and volatilities. The increase in likelihood from an unobserved component is large (about 65 points), and statistically significant at any reasonable confidence level. Based on recent data including the recent financial crisis, and a different modeling framework and estimation methodology, we confirm and extend the findings of Duffie et al. (2009) who point out the need for a latent component to prevent a downward bias in the estimation of default rate volatility and extreme default losses on portfolios of U.S. corporate debt. Our results indicate that the presence of a latent factor is not due to a few omitted macroeconomic covariates, but rather appears to capture different omitted effects at different times. In general, the default cycle and business cycle appear to depend on different processes. Inference on the default cycle using observed risk factors only is at best suboptimal, and at worst systematically misleading.

Third, we show that all three risk factors - common factors from observed macroeconomic and financial data, the latent frailty factor, and industry-specific contagion risk factors - are useful for out-of sample forecasting of default risk conditions. Feasible reductions in forecasting error are substantial, and far exceed the reductions achieved by standard models which use a limited set of observed covariates directly. Our findings lend support to models in which macroeconomic and default data are driven simultaneously by common factors. Our forecasting results do not lend support to models in which a few observed covariates drive defaults as exogenous factors directly. We find that forecasts improve most when an unobserved component is added to macro and contagion factors. Mean absolute forecasting errors reduce about 43% on average compared to a benchmark with observed risk factors only. Such reductions of more than 50% in most years are substantial and have clear practical implications for the computation of Value-at-Risk based capital buffers, for the stress testing of selected parts of the loan book, and the pricing of short-term debt. Reductions in MAE are most pronounced when frailty effects are highest. Examples are the year 2002, when default rates remain high while the economy is out of recession. Also, in the period 2005-07 leading up to the recent financial crisis, default conditions are substantially more benign than what is implied by observed macro data.

This paper proceeds as follows. In Section 2.2 we introduce the econometric framework which combines a nonlinear non-Gaussian panel time series model with an approximate dynamic factor model for many covariates. Section 2.3 shows how the proposed econometric model can be represented as a multi-factor firm value model for dependent defaults. In

Section 2.4 we discuss the estimation of the unknown parameters. Section 2.5 introduces the data for our empirical study, presents the major empirical findings, and discusses the out-of-sample forecasting results. Section 2.6 concludes.

2.2 The econometric framework

In this section we present our reduced form econometric model for dependent defaults. The economic implications of this framework are discussed in Section 2.3. We denote the default counts of cross section j at time t as y_{jt} for $j = 1, \dots, J$, and $t = 1, \dots, T$. The index j refers to a specific combination of firm characteristics, such as industry sector, current rating class, and company age. Defaults are correlated in the cross-section through exposure to the same business cycle, financing conditions, monetary and fiscal policy, firm and consumer sentiment, etcetera. The macroeconomic impact is summarized by exogenous factors in the $R \times 1$ vector F_t . Other explanatory covariates, such as trailing equity returns and volatilities, and trailing industry-level default rates, are collected in vector C_t . A frailty factor f_t^{uc} (where ‘uc’ refers to unobserved component) captures default clustering above and beyond what is implied by observed macro data. Subject to the conditioning on observed and unobserved risk factors, defaults occur independently in the cross section, see for example CreditMetrics (2007) or Lando (2003, Chapter 9). The panel time series of defaults is therefore modeled by

$$y_{jt}|F_t, C_t, f_t^{uc} \sim \text{Binomial}(k_{jt}, \pi_{jt}), \quad (2.1)$$

where y_{jt} is the total number of default ‘successes’ from k_{jt} exposures. Conditional on F_t , C_t and f_t^{uc} , the counts y_{jt} are assumed to be generated as independent Bernoulli-trials with time-varying default probability π_{jt} . In our model, k_{jt} represents the number of firms in cell j that are active at the beginning of period t . We recount exposures k_{jt} at the beginning of each quarter.

The measurement and forecasting of conditional default probability π_{jt} is our central focus. The probability π_{jt} can alternatively be referred to as hazard rates or default intensities in discrete time. We specify π_{jt} as the logistic transform of an index function θ_{jt} and therefore θ_{jt} can be interpreted as the log-odds or logit transform of π_{jt} . Probit and other transformations are also possible. Each specification implies a different model formulation and may lead to (slightly) different estimation results. We prefer the logit

transformation because of its simplicity. The default probabilities are specified by

$$\pi_{jt} = (1 + e^{-\theta_{jt}})^{-1}, \quad (2.2)$$

$$\theta_{jt} = \lambda_j + \beta_j f_t^{uc} + \gamma_j' F_t + \delta_j' C_t, \quad (2.3)$$

where λ_j is a fixed effect for the j th cross section. The coefficient vectors β_j , γ_j , and δ_j capture risk factor sensitivities, which may depend on firm characteristics such as industry sector and rating class. The time-varying default probabilities π_{jt} are determined by observed risk factors F_t and C_t as well as by the unobserved factor f_t^{uc} . The conditionally Binomial assumption for (2.1) is therefore analogous to the doubly-stochastic default times assumption of Azizpour and Giesecke (2008) and Duffie et al. (2009). The default signals θ_{jt} do not contain idiosyncratic error terms. Instead, idiosyncratic randomness is captured in (2.1). The log-odds of conditional default probabilities may vary over time due to variation in the macroeconomic factors, F_t , observed covariates, C_t , and the frailty component, f_t^{uc} .

The frailty factor f_t^{uc} is modeled by an unobserved dynamic process which we specify by the stationary autoregressive process of order one,

$$f_t^{uc} = \phi f_{t-1}^{uc} + \sqrt{1 - \phi^2} \eta_t, \quad \eta_t \sim \text{NID}(0, 1), \quad t = 1, \dots, T, \quad (2.4)$$

where $0 < \phi < 1$ and η_t is a serially uncorrelated sequence of standardized Gaussian disturbances. We therefore have $E(f_t^{uc}) = 0$, $\text{Var}(f_t^{uc}) = 1$, and $\text{Cov}(f_t^{uc}, f_{t-h}^{uc}) = \phi^h$. This specification enables the identification of β_j in (2.3). Extensions to multiple unobserved factors for firm-specific heterogeneity and to other dynamic specifications for f_t^{uc} are possible as is illustrated by Koopman and Lucas (2008).

Modeling the dependence of firm defaults on observed macro variables is an active area of current research, see Duffie, Saita, and Wang (2007), Duffie et al. (2009) and the references therein. The number of macroeconomic variables in the model differs across studies but is usually small. Instead of opting for a specific selection in our study, we collect a large number of macroeconomic and financial variables denoted by x_{nt} for $n = 1, \dots, N$. This time series panel of macroeconomic predictor variables typically contains many regressors. The panel is assumed to adhere to a factor structure as given by

$$x_{nt} = \Lambda_n F_t + \zeta_{nt}, \quad n = 1, \dots, N, \quad (2.5)$$

where F_t is a vector of principal components, Λ_n is a row vector of loadings, and ζ_{nt} is an idiosyncratic disturbance term. This static factor representation of the approximate

dynamic factor model (2.5) can be derived from a dynamic model specification, see Stock and Watson (2002a). This methodology of relating given variables of interest to a limited set of macroeconomic factors has been employed in the forecasting of inflation and production data, see Massimiliano, Stock, and Watson (2003), asset returns and volatilities, see Ludvigson and Ng (2007), and the term structure of interest rates, see Exterkate, van Dijk, Heij, and Groenen (2010). These studies have reported favorable results when such factors are used for forecasting.

The factors F_t can be estimated consistently using the method of principal components. This method is expedient for several reasons. First, dimensionality problems do not occur even for high values of N and T . This is particularly relevant for our empirical application, where $T, N > 100$ in both the macro and default datasets. Second, it can be shown that under relatively weak assumptions the method of principal components reduces to the maximum likelihood method when the idiosyncratic terms are assumed Gaussian. Third, the method can be extended to account for missing observations which are present in many macroeconomic time series panels. Finally, the extracted factors can be used for the forecasting of particular time series in the panel, see Forni, Hallin, Lippi, and Reichlin (2005). Equations (2.1) to (2.5) combine the approximate dynamic factor model with a non-Gaussian panel data model by inserting the elements of F_t from (2.5) into the signal equation (2.3). Parameter estimation is discussed in Section 2.4.

2.3 The financial framework

By relating the econometric model with the multi-factor model of CreditMetrics (2007) for dependent defaults, we can establish an economic interpretation of the parameters. In addition, we gain more intuition for the mechanisms of the model. Multi-factor models for firm default risk are widely used in risk management practice, see Lando (2003, Chapter 9).

In the special case of a standard static one-factor credit risk model for dependent defaults the values of the obligors' assets, V_i , are driven by a common random factor F , and an idiosyncratic disturbance ϵ_i . More specifically, the asset value of firm i , V_i , is modeled by

$$V_i = \sqrt{\rho_i} f + \sqrt{1 - \rho_i} \epsilon_i,$$

where scalar $0 < \rho_i < 1$ weights the dependence of firm i on the general economic condition factor f in relation to the idiosyncratic factor ϵ_i , for $i = 1, \dots, K$, where K is the number of firms, and where $(f, \epsilon_i)'$ has mean zero and variance matrix I_2 . The conditions in this

framework imply that

$$E(V_i) = 0, \quad \text{Var}(V_i) = 1, \quad \text{Cov}(V_i V_j) = \sqrt{\rho_i \rho_j},$$

for $i, j = 1, \dots, K$. In our multivariate dynamic model, the framework is extended into a more elaborate version for the asset value V_{it} of firm i at time t and is given by

$$\begin{aligned} V_{it} &= \omega_{i0} f_t^{uc} + \omega'_{i1} F_t + \omega'_{i2} C_t + \sqrt{1 - (\omega_{i0})^2 - \omega'_{i1} \omega_{i1} - \omega'_{i2} \omega_{i2}} \epsilon_{it} \\ &= \omega'_i \tilde{f}_t + \sqrt{1 - \omega'_i \omega_i} \epsilon_{it}, \quad t = 1, \dots, T, \end{aligned} \quad (2.6)$$

where frailty factor f_t^{uc} , macro factors F_t and firm/industry-specific covariates C_t have been introduced in (2.1), the associating weight vectors ω_{i0} , ω_{i1} , and ω_{i2} have appropriate dimensions, the factors and covariates are collected in $\tilde{f}_t = (f_t^{uc}, F'_t, C'_t)'$, and all weight vectors are collected in $\omega_i = (\omega_{i0}, \omega'_{i1}, \omega'_{i2})'$ with condition $\omega'_i \omega_i \leq 1$. The idiosyncratic standard normal disturbance ϵ_{it} is serially uncorrelated for $t = 1, \dots, T$. The unobserved component or frailty factor f_t^{uc} represents the credit cycle condition after controlling for the first M macro factors $F_{1,t}, \dots, F_{M,t}$ and the common variation in the covariates C_t . In other words, the frailty factor captures deviations of the default cycle from systematic macroeconomic and financial conditions. Without loss of generality we assume that all risk factors have zero mean and unit variance. Furthermore, we assume that the risk factors f_t^{uc} and F_t are uncorrelated with each other at all times.

In a firm value model, firm i defaults at time t when its asset value V_{it} drops below some threshold c_i , see Merton (1974) and Black and Cox (1976). In our framework, V_{it} is driven by systematic observed and unobserved factors as in (2.6). In our empirical specification, the threshold c_i depends on the current rating class, the industry sector, and the time elapsed since the initial rating assignment. For firms which have not defaulted yet, a default occurs when $V_{it} < c_i$ or, as implied by (2.6), when

$$\epsilon_{it} < \frac{c_i - \omega'_i \tilde{f}_t}{\sqrt{1 - \omega'_i \omega_i}}.$$

The conditional default probability is then given by

$$\pi_{it} = \Pr \left(\epsilon_{it} < \frac{c_i - \omega'_i \tilde{f}_t}{\sqrt{1 - \omega'_i \omega_i}} \right). \quad (2.7)$$

Favorable credit cycle conditions are associated with a high value of $\omega'_i \tilde{f}_t$ and therefore with a low default probability π_{it} for firm i . Furthermore, equation (2.7) can be related directly

to the econometric model specification in (2.2) and (2.3) where the firms ($i = 1, \dots, I$) are pooled into groups ($j = 1, \dots, J$) according to rating class, industry sector, and time from initial rating assignment. In particular, if ϵ_{it} is logistically distributed, we obtain

$$\begin{aligned} c_i &= \lambda_j \sqrt{1 - a_j}, & \omega_{i0} &= -\beta_j \sqrt{1 - a_j}, \\ \omega_{i1} &= -\gamma_j \sqrt{1 - a_j}, & \omega_{i2} &= -\delta_j \sqrt{1 - a_j}, \end{aligned}$$

where $a_j = (\beta_j^2 + \gamma_j' \gamma_j + \delta_j' \delta_j) / (1 + \beta_j^2 + \gamma_j' \gamma_j + \delta_j' \delta_j)$ for firm i that belongs to group j . The coefficient vectors λ_j , β_j , and γ_j are defined below (2.2) and (2.3). The parameters have therefore a direct interpretation in widely used portfolio credit risk models such as CreditMetrics (2007).

2.4 Estimation using state space methods

We next discuss parameter estimation and signal extraction of the factors for model (2.1) to (2.5). The estimation procedure for the macro factors is discussed in Section 2.4.1. The state space representation of the econometric model is provided in Section 2.4.2. We estimate the parameters using a computationally efficient procedure for Monte Carlo maximum likelihood and we extract the frailty factor from a similar Monte Carlo method. A brief outline of these procedures is given in Section 2.4.3. All computations are implemented using the Ox programming language and the associated set of state space routines from SsfPack, see Doornik (2007) and Koopman, Shephard, and Doornik (2008).

2.4.1 Estimation of the macro factors

The common factors F_t from the macro data are estimated by minimizing the objective function given by

$$\min_{\{F, \Lambda\}} V(F, \Lambda) = (NT)^{-1} \sum_{t=1}^T (X_t - \Lambda F_t)' (X_t - \Lambda F_t), \quad (2.8)$$

where the $N \times 1$ vector $X_t = (x_{1t}, \dots, x_{Nt})'$ contains macroeconomic variables and F is the set $F = \{F_1, \dots, F_T\}$ for the $R \times 1$ vector F_t . The observed stationary time series x_{nt} are demeaned and standardized to have unit unconditional variance for $n = 1, \dots, N$. Concentrating out F and rearranging terms shows that (2.8) is equivalent to maximizing $\text{tr}(\Lambda' S_{X'X} \Lambda)$ with respect to Λ and subject to $\Lambda' \Lambda = I_R$, where $S_{X'X} = T^{-1} \sum_t X_t X_t'$ is the sample covariance matrix of the data, see Lawley and Maxwell (1971) and Stock and Watson (2002a). The resulting principal components estimator of F_t is given by $\hat{F}_t = X_t' \hat{\Lambda}$,

where $\hat{\Lambda}$ collects the normalized eigenvectors associated with the R largest eigenvalues of $S_{X'X}$.

When the variables in X_t are not completely observed for $t = 1, \dots, T$, we employ the Expectation Maximization (EM) procedure as devised in the Appendix of Stock and Watson (2002b). This iterative procedure takes a simple form under the assumption that $x_{nt} \sim \text{NID}(\Lambda_n F_t, 1)$, where Λ_n denotes the n th row of Λ for $n = 1, \dots, N$. Here, $V(F, \Lambda)$ in (2.8) is a linear function of the log-likelihood $L(F, \Lambda | X^m)$ where X^m denotes the missing parts of the dataset X_1, \dots, X_T . Since $V(F, \Lambda)$ is proportional to $-L(F, \Lambda | X^m)$, the minimizers of $V(F, \Lambda)$ are also the maximizers of $L(F, \Lambda | X^m)$. This result is exploited in the EM algorithm of Stock and Watson (2002b) that we have adopted to compute \hat{F}_t for $t = 1, \dots, T$.

2.4.2 The factor model in state space form

We can formulate model (2.1) to (2.4) in state space form where F_t and C_t are treated as explanatory variables. In our implementation, F_t will be replaced by \hat{F}_t as obtained from the previous section. The estimation framework can therefore be characterized as a two-step procedure. By first estimating the principal components to summarize the variation in macroeconomic data, we have established a computationally feasible and relatively simple procedure. In Section 2.4.4 we present simulation evidence to illustrate the adequacy of our approach for parameter estimation and for uncovering the factors from the data.

The Binomial log-density function of model (2.1) is given by

$$\log p(y_{jt} | \pi_{jt}) = y_{jt} \log \left(\frac{\pi_{jt}}{1 - \pi_{jt}} \right) + k_{jt} \log(1 - \pi_{jt}) + \log \binom{k_{jt}}{y_{jt}}, \quad (2.9)$$

where y_{jt} is the number of defaults and k_{jt} is the number of firms in cross-section j , for $j = 1, \dots, J$ and $t = 1, \dots, T$. By substituting (2.2) for the default probability π_{jt} into (2.9) we obtain the log-density in terms of the log-odds ratio $\theta_{jt} = \log(\pi_{jt}) - \log(1 - \pi_{jt})$ given by

$$\log p(y_{jt} | \theta_{jt}) = y_{jt} \theta_{jt} + k_{jt} \log(1 + e^{\theta_{jt}}) + \log \binom{k_{jt}}{y_{jt}}. \quad (2.10)$$

The log-odds ratio is specified as

$$\theta_{jt} = Z_{jt} \alpha_t, \quad Z_{jt} = (e'_j, F'_t \otimes e'_j, C'_t \otimes e'_j, \beta_j), \quad (2.11)$$

where e_j denotes the j th column of the identity matrix of dimension J , the state vector

$\alpha_t = (\lambda_1, \dots, \lambda_J, \gamma_{1,1}, \dots, \gamma_{R,J}, \delta'_1, \dots, \delta'_J, f_t^{uc})'$ consists of the fixed effects λ_j together with the loadings $\gamma_{r,j}$ and δ'_j , and the unobserved component f_t^{uc} . The system vector Z_{jt} is time-varying due to the inclusion of F_t and C_t .

The state vector α_t contains all unknown coefficients that are linear in the signals θ_{jt} . The transition equation provides a model for the evolution of the state vector α_t over time and is given by

$$\alpha_{t+1} = T\alpha_t + Q\xi_t, \quad \eta_t \sim N(0, 1), \quad (2.12)$$

where the system matrices are given by

$$T = \text{diag}(I, \phi), \quad R = \begin{bmatrix} 0 \\ \sqrt{1 - \phi^2} \end{bmatrix},$$

and where η_t is the same as in (2.4). The initial elements of the state vector are subject to diffuse initial conditions except for f_t^{uc} , which has zero mean and unit variance.

The equations (2.10) and (2.12) belong to a class of non-Gaussian state space models as discussed in Durbin and Koopman (2001, Part II) and Koopman and Lucas (2008). In our formulation, most unknown coefficients are part of the state vector α_t and are estimated as part of the filtering and smoothing procedures described in Section 2.4.3. This formulation leads to a considerable increase in the computational efficiency of our estimation procedure. The remaining parameters are collected in a coefficient vector $\psi = (\phi, \beta_1, \dots, \beta_J)'$ and are estimated by the Monte Carlo maximum likelihood methods that we will discuss next.

2.4.3 Parameter estimation and signal extraction

Parameter estimation for a non-Gaussian model in state space form can be carried out by the method of Monte Carlo maximum likelihood. Once we have obtained an estimate of ψ , we can compute the conditional mean and variance estimates of the state vector α_t . In both cases we make use of importance sampling methods. The details of our implementation are given next.

For notational convenience we suppress the dependence of the density $p(y; \psi)$ on ψ . The likelihood function of our model (2.1) to (2.4) can be expressed by

$$\begin{aligned} p(y) &= \int p(y, \theta) d\theta = \int p(y|\theta) p(\theta) d\theta \\ &= \int p(y|\theta) \frac{p(\theta)}{g(\theta|y)} g(\theta|y) d\theta = E_g \left[p(y|\theta) \frac{p(\theta)}{g(\theta|y)} \right], \end{aligned} \quad (2.13)$$

where $y = (y_{11}, y_{21}, \dots, y_{JT})'$, $\theta = (\theta_{11}, \theta_{21}, \dots, \theta_{JT})'$, $p(\cdot)$ is a density function, $p(\cdot, \cdot)$ is a joint density, $p(\cdot|\cdot)$ is a conditional density, $g(\theta|y)$ is a Gaussian importance density, and E_g denotes expectations with respect to $g(\theta|y)$. The importance density $g(\theta|y)$ is constructed as the Laplace approximation to the intractable density $p(\theta|y)$. Both densities have the same mode and curvature at the mode, see Durbin and Koopman (2001) for details. Conditional on θ , we can evaluate $p(y|\theta)$ by

$$p(y|\theta) = \prod_{j,t} p(y_{jt}|\theta_{jt}).$$

It follows from (2.3) and (2.4) that the marginal density $p(\theta)$ is Gaussian and therefore $p(\theta) = g(\theta)$. Since $g(\theta|y)g(y) \equiv g(y|\theta)g(\theta)$ we obtain

$$p(y) = E_g \left[p(y|\theta) \frac{p(\theta)}{g(y|\theta)} \frac{g(y)}{p(\theta)} \right] = E_g \left[g(y) \frac{p(y|\theta)}{g(y|\theta)} \right] = g(y) E_g [w(y, \theta)], \quad (2.14)$$

where $w(y, \theta) = p(y|\theta)/g(y|\theta)$. A Monte Carlo estimator of $p(y)$ is therefore given by

$$\hat{p}(y) = g(y)\bar{w},$$

with

$$\bar{w} = M^{-1} \sum_{m=1}^M w^m = M^{-1} \sum_{m=1}^M \frac{p(y|\theta^m)}{g(y|\theta^m)},$$

where $w^m = w(\theta^m, y)$ is the value of the importance weight associated with the m -th draw θ^m from $g(\theta|y)$, and M is the number of Monte Carlo draws. The Gaussian importance density $g(\theta|y)$ is chosen for convenience and since it is possible to generate a large number of draws θ^m from it in a computationally efficient manner using the simulation smoothing algorithms of de Jong and Shephard (1995) and Durbin and Koopman (2002). We estimate the log-likelihood as $\log \hat{p}(y) = \log \hat{g}(y) + \log \bar{w}$, and include a bias correction term as discussed in Durbin and Koopman (1997).

The Gaussian importance density $g(\theta|y)$ is based on the approximating Gaussian model as given by

$$y_{jt} = c_{jt} + \theta_{jt} + u_{jt}, \quad u_{jt} \sim \text{NID}(0, d_{jt}), \quad (2.15)$$

where the disturbances u_{jt} are mutually and serially uncorrelated, for $j = 1, \dots, j$ and $t = 1, \dots, T$. The unknown constant c_{jt} and variance d_{jt} are determined by the individual matching of the first and second derivative of $\log p(y_{jt}|\theta_{jt})$ in (2.10) and $\log g(y_{jt}|\theta_{jt}) = -\frac{1}{2} \log 2\pi - \frac{1}{2} \log d_{jt} - \frac{1}{2} d_{jt}^{-1} (y_{jt} - c_{jt} - \theta_{jt})^2$ with respect to the signal θ_{jt} . The matching equations for c_{jt} and d_{jt} rely on θ_{jt} for each j, t . For an initial value of θ_{jt} , we compute

c_{jt} and d_{jt} for all j, t . The Kalman filter and smoother compute the estimates for signal θ_{jt} based on the linear Gaussian state space model (2.15), (3.5) and (2.12). We compute new values for c_{jt} and d_{jt} based on the new signal estimates of θ_{jt} . We can repeat the computations for each new estimate of θ_{jt} . The iterations proceed until convergence is achieved, that is when the estimates of θ_{jt} do not change. The number of iterations for convergence are usually as low as 5 to 10 iterations. When convergence has taken place, the Kalman filter and smoother applied to the approximating model (2.15) compute the mode estimate of $\log p(\theta|y)$; see Durbin and Koopman (1997) for further details. A new approximating model needs to be constructed for each log-likelihood evaluation when the value for parameter vector ψ has changed. Finally, standard errors for the parameters in ψ are constructed from the numerical second derivatives of the log-likelihood function, that is

$$\hat{\Sigma} = \left[-\frac{\partial^2 \log p(y)}{\partial \psi \partial \psi'} \Big|_{\psi=\hat{\psi}} \right]^{-1}.$$

For the estimation of the latent factor f_t^{uc} and fixed coefficients in the state vector, we estimate the conditional mean of α by

$$\begin{aligned} \bar{\alpha} &= \text{E}[\alpha|y] = \int \alpha p(\alpha|y) d\alpha \\ &= \int \alpha \frac{p(\alpha|y)}{g(\alpha|y)} g(\alpha|y) d\alpha = \text{E}_g \left[\alpha \frac{p(\alpha|y)}{g(\alpha|y)} \right]. \end{aligned}$$

In a similar way as the development in (2.14), we obtain

$$\bar{\alpha} = \frac{\text{E}_g[\alpha w(\theta, y)]}{\text{E}_g[w(\theta, y)]},$$

since $p(\alpha) = g(\alpha)$, $p(y|\alpha) = p(y|\theta)$ and $g(y|\alpha) = g(y|\theta)$. The Monte Carlo estimator for $\bar{\alpha}$ is then given by

$$\hat{\alpha} = \widehat{\text{E}}[\alpha|y] = \left[\sum_{m=1}^M w^m \right]^{-1} \sum_{m=1}^M \alpha^m w^m,$$

where $\alpha^m = (\alpha_{11}^m, \dots, \alpha_{JT}^m)'$ is the m -th draw from $g(\alpha|y)$ and where θ^m is computed using (3.5), that is $\theta_{jt}^m = Z_{jt} \alpha_{jt}^m$ for $j = 1, \dots, J$ and $t = 1, \dots, T$. The associated conditional variances are given by

$$\widehat{\text{Var}}[\alpha_{jt}|y] = \left(\left[\sum_{m=1}^M w^m \right]^{-1} \sum_{m=1}^M (\alpha_{jt}^m)^2 w^m \right) - (\hat{\alpha}_{jt})^2,$$

and allow the construction of standard error bands.

In our empirical study we also present mode estimates for signal extraction and out-of-sample forecasting of default probabilities or hazard rates in (2.3). The mode estimates of α_{jt} are obtained by the Kalman filter smoother applied to the state space model (2.15), (3.5) and (2.12) where c_{jt} and d_{jt} are computed by using the mode estimate of θ_{jt} . Finally, the mode estimate of $\pi = \pi(\theta)$ is given by $\bar{\pi} = \pi(\bar{\theta})$ for any nonlinear function $\pi(\cdot)$ that is known and has continuous support. We refer to Durbin and Koopman (2001, Chapter 11) for further details.

2.4.4 Simulation experiments

In this subsection we investigate whether the econometric methods of Sections 2.4.1 and 2.4.3 can distinguish the variation in default conditions due to changes in the macroeconomic environment from changes in unobserved frailty risk. The first source is captured by principal components F_t , while the second source is estimated via the unobserved factor f_t^{uc} . This exercise is important since estimation by Monte Carlo maximum likelihood should not be biased towards attributing variation to a latent component when it is due to an exogenous covariate. For this purpose we carry out a simulation study that is close to our empirical application in Section 2.5. The variables are generated by the equations

$$\begin{aligned} F_t &= \Phi_F F_{t-1} + u_{F,t}, & u_{F,t} &\sim N(0, I - \Phi_F \Phi_F'), \\ e_t &= \Phi_I e_{t-1} + u_{I,t}, & u_{I,t} &\sim N(0, I - \Phi_I \Phi_I'), \\ X_t &= \Lambda F_t + e_t, \\ f_t^{uc} &= \phi_{uc} f_{t-1}^{uc} + u_{f,t}, & u_{f,t} &\sim N(0, 1 - \phi_{uc}^2), \end{aligned}$$

where ϕ_{uc} and the elements of the matrices Φ_F , Φ_I , and Λ are generated for each simulated dataset from the uniform distribution $U[.,.]$, that is $\phi_{uc} \sim U[0.6, 0.8]$, $\Phi_F(i, j) \sim U[0.6, 0.8]$, $\Phi_I(i, j) \sim U[0.2, 0.4]$, and $\Lambda(i, j) \sim U[0, 2]$, where $A(i, j)$ is the (i, j) th element of matrix $A = \Phi_F, \Phi_I, \Lambda$. For computational convenience we consider F_t to be a scalar process ($M = 1$) and we have no firm-specific covariates ($C_t = 0$). The default counts y_{jt} in pooling group j are generated by the equations

$$\begin{aligned} \theta_{jt} &= \lambda_j + \beta f_t^{uc} + \gamma F_t, \\ y_{jt} &\sim \text{Binomial}(k_{jt}, (1 + \exp[-\theta_{jt}])^{-1}), \end{aligned}$$

where f_t^{uc} and F_t represent their simulated values, and exposure counts k_{jt} come from the dataset which is explored in the next section. The parameters λ_j , β , γ are chosen similar to their maximum likelihood values reported in Section 2.5. Simulation results are

based on 1000 simulations. Each simulation uses 50 importance samples during simulated maximum likelihood estimation, and 500 importance samples for signal extraction.

A selection of the graphical output from our Monte Carlo study is presented in Figure 2.1. We find that the principal components estimate \hat{F} captures the factor space F well. The goodness-of-fit statistic R^2 is on average 0.94. The conditional mean estimate of f^{uc} is close to the simulated unobserved factor, with an average R^2 of 0.73. The sampling distributions of ϕ_{uc} and λ_0 appear roughly symmetric and Gaussian, while the distributions of factor sensitivities β_0 and γ_1 appear skewed to the right. This is consistent with their interpretation as factor standard deviations. The distributions of ϕ_{uc} , β_0 , λ_0 , and γ_1 are all centered around their true values. We conclude that our modeling framework enables us to discriminate between possible sources of default rate variation. The resulting parameter estimates are overall correct for both ψ and state vector α .

Finally, the standard errors for the estimated factor loadings γ do not take into account that the principal components are estimated with some error in a first step. We therefore need to investigate whether this impairs inference on these factor loadings. In each simulation we estimate parameters and associated standard errors using true factors F_t as well as their principal components estimates \hat{F}_t . The bottom panel in Figure 2.1 plots the empirical distribution functions of t-statistics associated with testing the null hypothesis $H_0 : \gamma_1 = 0$ when either F_t or \hat{F}_t is used. The t-statistics are very similar in both cases. Other standard errors are similarly unaffected. We conclude that the substitution of F_t with \hat{F}_t has negligible effects for parameter estimation.

2.5 Estimation results and forecasting accuracy

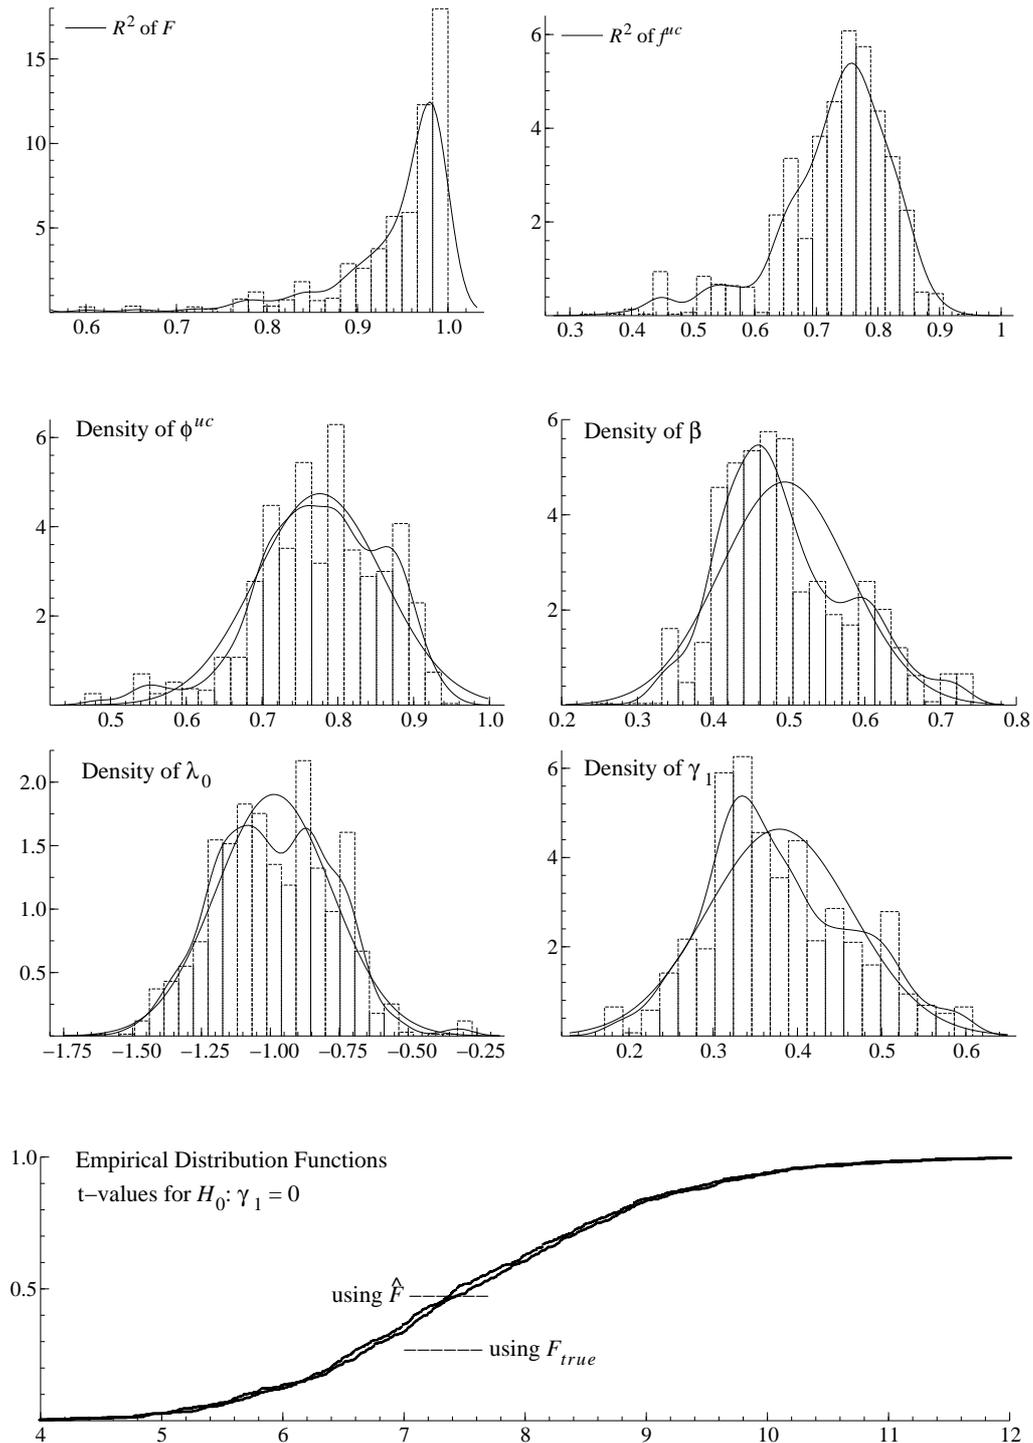
We first describe the macroeconomic, financial, and firm default data used in our empirical study. We then discuss our main findings from the study. We conclude with the discussion of out-of-sample forecasting results for a cross-section of default hazard rates.

2.5.1 Data

We use data from two main sources. First, a panel of more than 100 macroeconomic and financial time series is constructed from the Federal Reserve Economic Database FRED (<http://research.stlouisfed.org/fred2>). The aim is to select series which contain information about systematic credit risk conditions. The variables are grouped into five broad categories: (a) bank lending conditions, (b) macroeconomic and business cycle indicators, including labor market conditions and monetary policy indicators, (c)

Figure 2.1: Simulation analysis

Graphs 1 and 2 contain the sampling distributions of R-squared goodness-of-fit statistics in regressions of \hat{F} on simulated factors F , and conditional mean estimates $\hat{E}[f^{uc}|y]$ on the true f^{uc} , respectively. Graphs 3 to 6 plot the sampling distributions of key parameters ϕ_{uc} , β , λ_0 , and γ_1 . The bottom panel plots two empirical distribution functions of the t-statistics associated with testing $H_0 : \gamma_1 = 0$. In each simulation either F or \hat{F} are used to obtain Monte Carlo maximum likelihood parameter and standard error estimates. Distribution plots are based on 1000 simulations. The dimensions of the default panel are $N=112$, and $T=100$. The macro panel has $N=120$, and $T=100$.



open economy macroeconomic indicators, (d) micro-level business conditions such as wage rates, cost of capital, and cost of resources, and (e) stock market returns and volatilities. The macro variables are quarterly time series from 1970Q1 to 2009Q4. Table 2.1 presents a listing of the series for each category. The macroeconomic panel contains both current information indicators (real GDP, industrial production, unemployment rate) and forward looking variables (stock prices, interest rates, credit spreads, commodity prices).

A second dataset is constructed from the default data of Moody's. The database contains rating transition histories and default dates for all rated firms from 1981Q1 to 2009Q4. This data contains the information to determine quarterly values for y_{jt} and k_{jt} in (2.1). The database distinguishes 12 industries which we pool into $D = 7$ industry groups: banks and financials (fin); transport and aviation (tra); hotels, leisure, and media (lei); utilities and energy (egy); industrials (ind); technology and telecom (tec); retailing and consumer goods (rcg). We further consider four age cohorts: less than 3, 3 to 6, 6 to 12, and more than 12 years from the time of the initial rating assignment. Age cohorts are included since default probabilities may depend on the age of a company. A proxy for age is the time since the initial rating has been established. Finally, there are four rating groups, an investment grade group $Aaa - Baa$, and three speculative grade groups Ba , B , and $Caa - C$. Pooling over investment grade firms is necessary since defaults are rare in this segment. In total we distinguish $J = 7 \times 4 \times 4 = 112$ different groups.

In the process of counting exposures and defaults, a previous rating withdrawal is ignored if it is followed by a later default. If there are multiple defaults per firm, we consider only the first event. In addition, we exclude defaults that are due to a parent-subsidiary relationship. Such defaults typically share the same default date, resolution date, and legal bankruptcy date in the database. Inspection of the default history and parent number confirms the exclusion of these cases.

Aggregate default counts, exposure counts, and fractions are presented in the top panel of Figure 2.2. We observe pronounced default clustering around the recession years of 1991, 2001, and the recent financial crisis of 2007-09. Since defaults cluster due to high levels of latent systematic risk, it follows that systematic risk is serially correlated and may also account for the autocorrelation in aggregate defaults. Defaults may already rise before the onset of a recession, for example, in the years 1990 and 2000, and they may remain elevated as the economy recovers from recession, for example, in the year 2002. The bottom panel of Figure 2.2 presents disaggregated default fractions for four broad rating groups. Default clustering is visible for all rating groups.

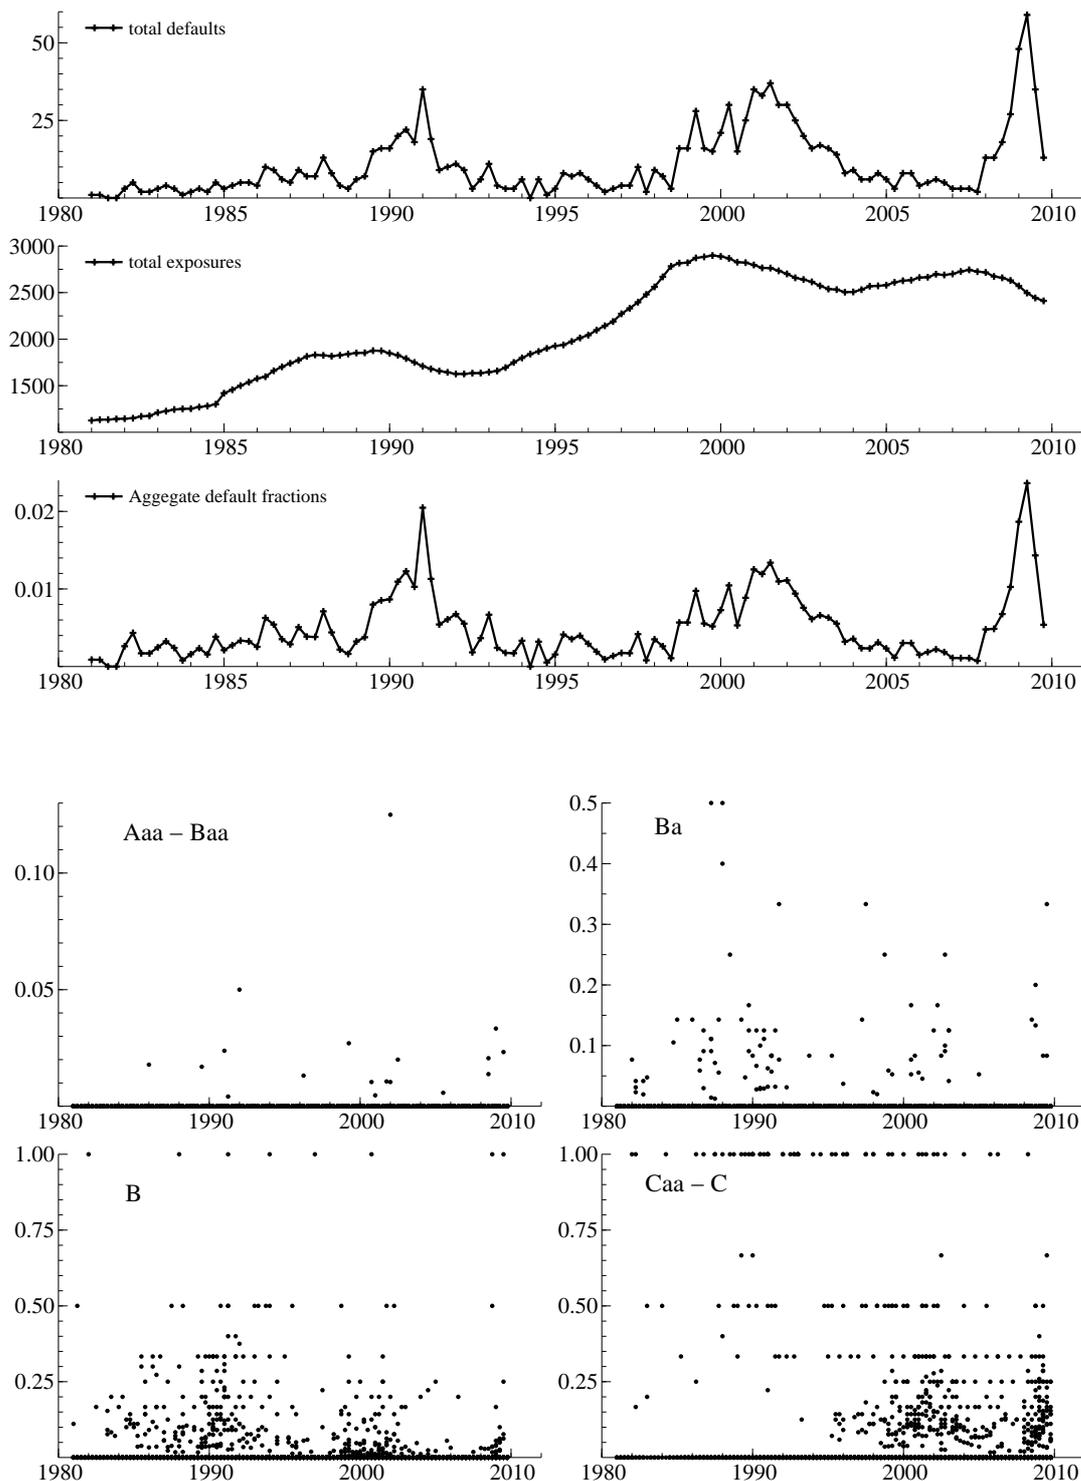
Our proposed model considers groups of firms rather than individual firms. As a result it is not straightforward to include firm specific information beyond rating classes

Table 2.1: Macroeconomic and financial predictor variables

Main category	Summary listing	Total
(a) Bank lending conditions		
Size of overall lending	Total Commercial Loans Total Real Estate Loans Total Consumer Credit outst. Commercial&Industrial Loans Bank loans and investments Household obligations/income	Household debt/income-ratio Federal debt of Non-fin. sector Excess Reserves of Dep. Institutions Total Borrowings from Fed Reserve Household debt service payments Total Loans and Leases, all banks
Extend of problematic banking business	Non-performing Loans Ratio Net Loan Losses Return on Bank Equity Non-perf. Commercial Loans	Non-performing Total Loans Total Net Loan Charge-offs Loan Loss Reserves
(b) Macro and BC conditions		
General macro indicators	Real GDP Industr. Production Index Private Fixed Investments National Income Manuf. Sector Output Manuf. Sector Productivity Government Expenditure	ISM Manufacturing Index Uni Michigan Consumer Sentiment Real Disposable Personal Income Personal Income Consumption Expenditure Expenditure Durable Goods Gross Private Domestic Investment
Labour market conditions	Unemployment rate Weekly hours worked Employment/Population-Ratio	Total No Unemployed Civilian Employment Unemployed, more than 15 weeks
Business Cycle leading/ coinciding indicators	New Orders: Durable goods New orders: Capital goods Capacity Util. Manufacturing Capacity Util. Total Industry Light weight vehicle sales Housing Starts New Building Permits	Final Sales of Dom. Product Inventory/Sales-ratio Change in Private Inventories Inventories: Total Business Non-farm housing starts New houses sold Final Sales to Domestic Buyers
Monetary policy indicators	M2 Money Stock UMich Infl. Expectations Personal Savings Gross Saving	CPI: All Items Less Food CPI: Energy Index Personal Savings Rate GDP Deflator, implicit
Firm Profitability	Corp. Profits Net Corporate Dividends	After Tax Earnings Corporate Net Cash Flow
(c) Intern'l competitiveness		
Terms of Trade	Trade Weighted USD	FX index major trading partners
Balance of Payments	Current Account Balance Balance on Merchandise Trade	Real Exports Goods, Services Real Imports Goods & Services
(d) Micro-level conditions		
Labour cost/wages	Unit Labor Cost: Manufacturing Total Wages & Salaries Management Salaries Technical Services Wages Employee Compensation Index	Unit Labor Cost: Nonfarm Business Non-Durable Manufacturing Wages Durable Manufacturing Wages Employment Cost Index: Benefits Employment Cost Index: Wages & Salaries
Cost of capital	1Month Commerical Paper Rate 3Month Commerical Paper Rate Effective Federal Funds Rate AAA Corporate Bond Yield BAA Corporate Bond yield	Treasury Bond Yield, 10 years Term Structure Spread Corporate Yield Spread 30 year Mortgage Rate Bank Prime Loan Rate
Cost of resources	PPI All Commodities PPI Interm. Energy Goods PPI Finished Goods	PPI Industrial Commodities PPI Crude Energy Materials PPI Intermediate materials
(e) Equity market conditions		
Equity Indexes and respective volatilities	S&P 500 Nasdaq 100 S&P Small Cap Index	Dow Jones Industrial Average Russell 2000

Figure 2.2: Aggregated default data and disaggregated fractions

The top graph presents time series plots of (a) the total default counts $\sum_j y_{jt}$ aggregated to a univariate series, (b) total number of firms $\sum_j k_{jt}$ in the database, and (c) aggregate default fractions $\sum_j y_{jt} / \sum_j k_{jt}$ over time. The bottom graph plots disaggregated default fractions y_{jt}/k_{jt} over time for four broad rating groups *Aaa – Baa*, *Ba*, *B*, and *Caa – C*. Each plot contains multiple default fractions over time, disaggregated across industries and time from initial rating assignment.



and industry sectors. Firm-specific covariates such as equity returns, volatilities and leverage are found to be important in Vassalou and Xing (2004), Duffie et al. (2007), and Duffie et al. (2009). We acknowledge that ratings alone are unlikely to be sufficient statistics for future default. To accommodate this concern to some extent, the set of covariates in the model is extended with average measures of firm-specific variables across firms in the same industry groups. We use the S&P industry-level equity index data from Datastream to construct trailing equity return and spot volatility measures at the industry level. The equity volatilities at the industry level are constructed as realized variance estimates based on average squared monthly returns over the past year. We also follow Das, Duffie, Kapadia, and Saita (2007) and Duffie et al. (2009) by including the trailing 1-year return of the S&P 500 stock index, an S&P 500 spot volatility measure, and the 3-month T-bill rate from Datastream. These additional observed risk factors are treated in the same way as the first 10 principal components from the macroeconomics dataset.

2.5.2 Macro and contagion factors

In Figure 2.3 we present the ten principal components obtained from the macro panel of Table 2.1 and computed by the EM procedure of Section 2.4.1. The NBER recession dates are depicted as shaded areas. The estimated first factor from the macroeconomic and financial panel is mainly associated with production and employment data; it accounts for a large share of 24% of total variation in the panel. The first factor exhibits clear peaks around the U.S. business cycle troughs. The remaining factors also have peaks and troughs around these periods, but the association with the U.S. business cycle is less strong. Overall, we select $M = 10$ factors which capture 82% of the variation in the panel.

Default contagion is a possible alternative source of default clustering in observed data, see Jorion and Zhang (2007b), Lando and Nielsen (2008), and Boissay and Gropp (2010). We assume that default contagion due to supply chain relationships is most important at the intra-industry level. For example, a defaulting manufacturing firm may weaken other up- or downstream manufacturing firms. Similarly, a defaulting financial firm is assumed to affect other financial firms. To capture industry-level (contagion) dynamics, we regress trailing one year default rates at the industry-level on a constant and the trailing one year aggregate default rate. Contagion factors are then obtained as the resulting standardized residuals. In this way, we eliminate the effect of the common factors F_t and f_t^{uc} and we retain industry-specific variation.

Figure 2.4 presents our estimated contagion factors for seven broad industry groups.

Figure 2.3: Principal components from unbalanced macro data

The figure plots the first ten principal components from unbalanced macro and financial time series data as listed in Table 2.1. Shaded areas indicate NBER recession periods.

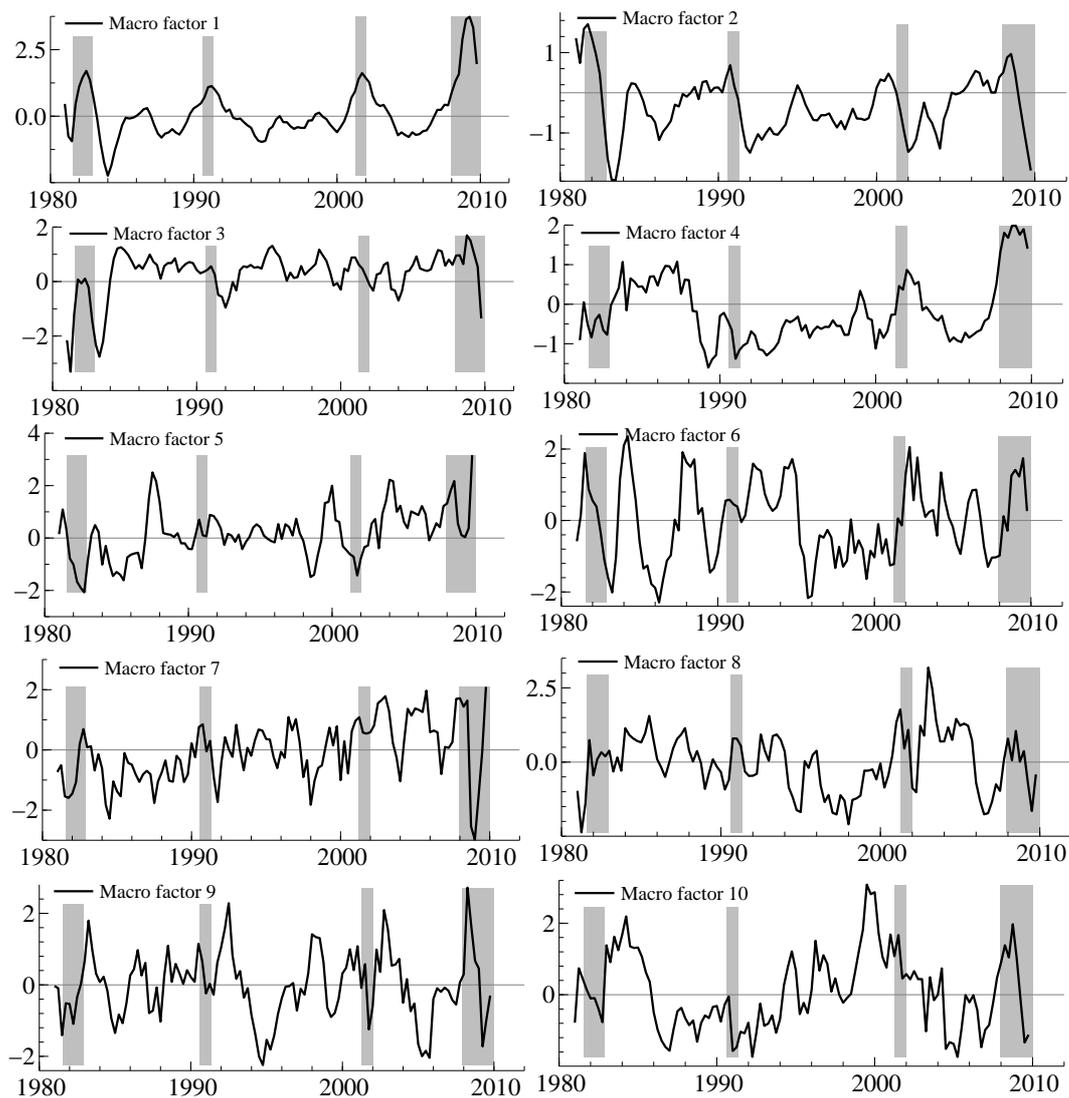
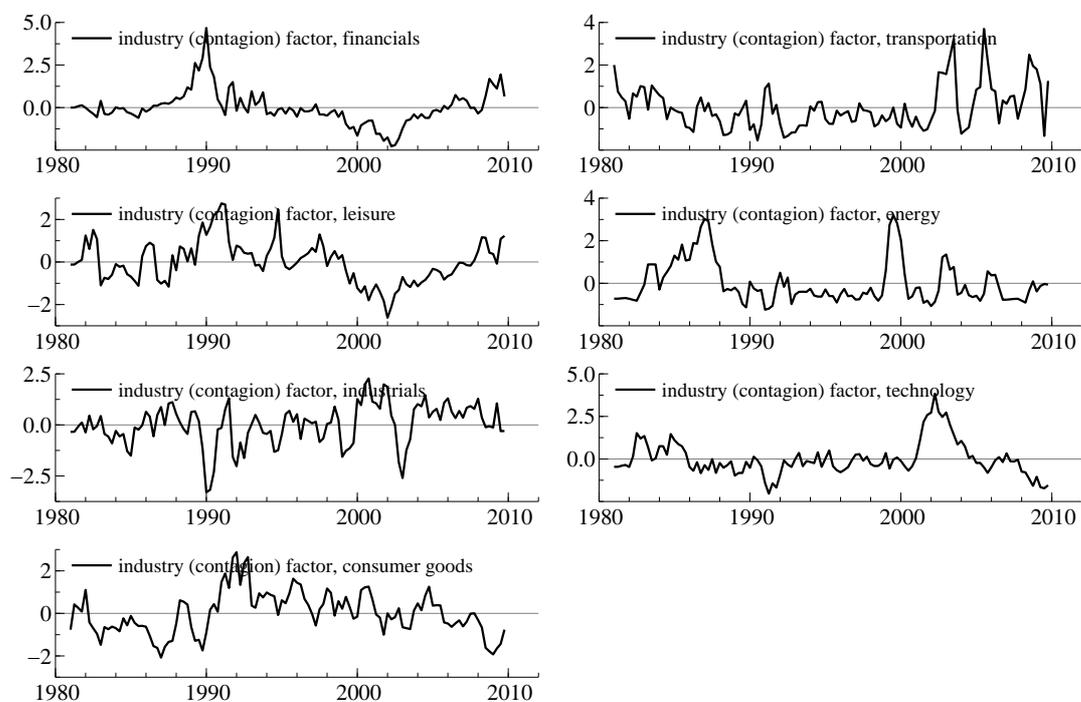


Figure 2.4: Industry-specific contagion factors

We plot observed industry-specific contagion risk factors for seven industries. The factors are obtained by regression of trailing one-year industry-level default rates on a constant and the trailing total default rate. Factors are standardized to unit variance.



For financial firms, we observe the savings and loans crisis of the late 1980s, the relatively mild impact of the 2001 recession on financials and the financial crisis in 2008-2009. In other sectors, we observe the effects of the the burst of the dot-com bubble on technology firms in 2001-2002, and the effects of the 9/11 attacks on the US transportation and aviation sector in 2002. A contagion interpretation may be appropriate in some cases. We conclude that the contagion factors capture the salient features in defaults at the industry level.

2.5.3 Model specification

The model specification for the default counts of our $J = 112$ groups is as follows. The individual time series of counts is modelled as a Binomial sequence with log-odds ratio θ_{jt} as given by (2.3) or (3.5) where the scalar coefficient λ_j is a fixed effect, scalar β_j pertains to the frailty factor, vector γ_j to the principal components and vector δ_j to the contagion factors, for $j = 1, \dots, J$. The model includes ten principal components that capture 82% of the variation from 107 macroeconomic and financial predictor variables, equity returns and volatilities at the industry level, industry-specific contagion factors, and the firm-specific ratings, industry group, and age cohorts.

Since the cross-section is high-dimensional, we follow Koopman and Lucas (2008) in reducing the number of parameters by restricting the coefficients in the following additive structure

$$\bar{\chi}_j = \chi_0 + \chi_{1,d_j} + \chi_{2,a_j} + \chi_{3,s_j}, \quad \bar{\chi} = \lambda, \beta, \gamma, \delta, \quad (2.16)$$

where χ_0 represents the baseline effect, $\chi_{1,d}$ is the industry-specific deviation, $\chi_{2,a}$ is the deviation related to age and $\chi_{3,s}$ is the deviation related to rating group. The deviations of all seven industry groups (fin, tra, lei, egy, tec, ind, and rcg) cannot be identified simultaneously given the presence of χ_0 . To identify the model, we assume that $\chi_{1,d_j} = 0$ for the retail and consumer goods group, $\chi_{2,a_j} = 0$ for the age group of 12 years or more, and $\chi_{3,s_j} = 0$ for the rating rating group $Caa - C$. These normalizations are innocuous and can be replaced by alternative baseline choices without affecting our conclusions. For the frailty factor coefficients, we do not account for age and therefore set $\beta_{2,a} = 0$ for all a . For the principal components coefficients, we only account for rating groups and therefore we have $\gamma_{1,d} = 0$ and $\gamma_{2,a} = 0$, for all d, s . For the contagion factor coefficients, we only account for industry groups and therefore we have $\delta_{2,a} = 0$ and $\delta_{3,s} = 0$, for all d, s . Using this parameter specification, we combine model parsimony with the ability to test a rich set of hypotheses empirically given the data at hand.

2.5.4 Empirical findings

Table 2.2 presents the parameter estimates for three different specifications of the signal equation (2.3). Model 1 does not contain the macro factors, $\beta_j = 0$. Model 2 does not contain the latent risk factors, $\gamma_{rj} = 0$ for all r and j . Model 3 refers to specification (2.3) without restrictions.

When comparing the log-likelihood values of Models 1 and 3, we can conclude that adding a latent dynamic frailty factor increases the log-likelihood by approximately 65 points. This increase is statistically significant at the 1% level. Since in practice most default models rely on a set of covariates, this finding indicates that a model without a frailty factor can systematically provide misleading indications of default conditions. Therefore, the industry practise is at best suboptimal, and at worst systematically misleading when used for inference on default conditions. Furthermore, our finding supports Duffie et al. (2009), who argue that firms are exposed to a common dynamic latent component driving default in addition to observed risk factors. Ignoring this component leads to a significant downward omitted variable bias when assessing the default rate volatility and the probability of extreme default losses.

We further find that Model 2 produces a better in-sample fit to the data than Model 1 in terms of the maximized log-likelihood value. Hence, a single unobserved component captures default conditions better than ten principal components from the macroeconomic panel. We therefore conclude that business cycle dynamics and default risk conditions are different processes. This finding is relevant for credit risk managers in financial institutions and for policy makers in charge of financial stability.

The principal components also capture covariation in defaults. The difference in the log-likelihood values of Models 2 and 3 is 44 points and is significant at a 5% level. We may therefore conclude that all risk factors in our model are significant. However, all principal components are not of equal importance to default rates. For example, factors 3 and 6 capture 10% and 4% of the variation in the macro panel, respectively, but they have no effect on default counts.

The industry-specific contagion factors are significant for explaining defaults. For financial firms, loadings on the contagion factors are estimated as positive values and significant. For very competitive industries such as transportation and aviation, we obtain negative loadings with respect to trailing industry-level default rates. It may indicate that competitive effects from trailing defaults offset contagion effects in some industries. Overall we can conclude that trailing one-year industry-level default rates are good predictors of future default rates in specific industries.

Table 2.2: Estimation results

We report the maximum likelihood estimates of selected coefficients in the specification for the signal or log-odds ratio (2.3) with parameterization $\bar{\chi}_j = \chi_0 + \chi_{1,d_j} + \chi_{2,a_j} + \chi_{3,s_j}$ for $\bar{\chi} = \lambda, \beta$. Coefficients λ refer to fixed effects or baseline hazard, coefficients β refer to the frailty factor, and coefficients γ and δ refer to the macro and contagion factors, respectively. Monte Carlo log-likelihood evaluation is based on $M = 5000$ importance samples. Data is from 1981Q1 to 2009Q4. Further details of the model specification are discussed in Section 2.5.4.

par	Model 1: Only F_t		Model 2: Only f_t^{uc}		Model 3: All Factors	
	val	t-val	val	t-val	val	t-val
λ_0	-2.62	12.72	-2.56	9.19	-2.88	10.24
$\lambda_{1,fin}$	0.01	0.10	0.06	0.45	0.03	0.19
$\lambda_{1,tra}$	0.19	1.13	0.18	1.16	0.19	1.24
$\lambda_{1,lei}$	0.00	0.00	-0.09	0.78	-0.04	0.31
$\lambda_{1,egy}$	-0.21	1.60	-0.05	0.25	-0.44	2.15
$\lambda_{1,ind}$	-0.11	1.28	-0.19	1.85	-0.12	1.19
$\lambda_{1,tec}$	-0.28	2.04	-0.25	2.10	-0.29	2.19
$\lambda_{2,0-3}$	-0.20	1.74	-0.23	2.06	-0.25	2.21
$\lambda_{2,4-5}$	0.26	2.39	0.17	1.46	0.14	1.38
$\lambda_{2,6-12}$	0.24	2.04	0.14	1.20	0.15	1.30
$\lambda_{3,IG}$	-7.55	14.46	-6.98	16.80	-7.56	13.69
$\lambda_{3,Ba}$	-3.26	12.07	-3.21	15.62	-3.88	13.68
$\lambda_{3,B}$	-1.21	6.51	-1.25	7.78	-1.79	7.86
β_0			0.60	4.90	0.53	4.30
$\beta_{1,fin}$			-0.14	0.81	-0.18	1.45
$\beta_{1,tra}$			0.03	0.15	0.04	0.30
$\beta_{1,lei}$			0.13	1.04	-0.00	0.03
$\beta_{1,egy}$			-0.37	1.92	0.52	2.80
$\beta_{1,ind}$			0.15	1.20	-0.17	1.93
$\beta_{1,tec}$			-0.02	0.15	-0.07	0.66
$\beta_{2,IG}$			0.36	1.18	0.07	0.18
$\beta_{2,Ba}$			0.23	1.38	0.44	1.92
$\beta_{2,B}$			0.20	2.22	0.35	2.54
γ_1^{IG}	1.37	4.02			1.44	4.47
γ_1^{Ba}	0.49	2.19			0.68	3.18
γ_1^B	0.44	5.01			0.63	3.86
γ_1^{Caa}	0.42	3.62			0.53	4.02

δ_{fin}	0.18	2.33	0.17	2.31	0.17	2.00
δ_{tra}	-0.45	3.57	-0.37	2.72	-0.34	2.76
δ_{lei}	0.04	0.52	0.13	1.90	0.11	1.71
δ_{egy}	0.27	3.04	0.19	2.22	-0.03	0.34
δ_{ind}	0.02	0.29	-0.02	0.42	-0.05	1.13
δ_{tec}	0.30	5.21	0.25	3.98	0.25	3.53
δ_{rcg}	0.19	2.78	0.19	2.71	0.17	2.54
LogLik	-2660.98		-2639.43		-2595.87	

2.5.5 Interpretation of the frailty factor

We have given evidence in Section 2.5.4 that firms are exposed to a common dynamic latent factor driving default after controlling for measurable risk factors. Given its statistical and economic significance, we may conclude that the business cycle and the default cycle are related but depend on different processes. The approximation of the default cycle by business cycle indicators may not be sufficiently accurate. Figure 2.5 presents the frailty factor estimates for Models 2 and 3. The recession periods of 1983, 1991, 2001, and 2008-09 are marked as shaded areas. Recession periods coincide with peaks in the default cycle in the top panel for Model 2. The bottom panel presents the estimated frailty effects for Model 3.

Duffie et al. (2009) suggest that the frailty factor captures omitted relevant covariates together with other omitted effects that are difficult to quantify. Our results suggest that the frailty factor captures only other omitted effects which can be different at different times. The frailty effects in the period 2001-2002 can be attributed to the disappearance of trust in the accuracy of public accounting information following the Enron and WorldCom scandals. While the effects are important for accessing credit, they are difficult to quantify. Similarly, the downward movements of the frailty factor in 2005-2007 suggest that Model 3 is able to capture the positive effects of recent advances in credit risk transfer and securitization. These advances have led to cheap credit access. The estimated frailty factor appears to capture different omitted effects at different times, rather than that it substitutes for a single missing covariate.

Figure 2.6 presents the estimated composite default signals θ_{jt} for investment grade firms (Aaa-Baa) against low speculative grade firms (Caa-C). The frailty effects are less important for investment grade firms. The default clustering implied by observed risk factors is sufficient to match the default intensities in the recession periods 1983, 1991, 2001, and 2008. For the low speculative grade group, frailty effects indicate additional default clustering in the 1980s, and also during the 1991 recession. The bottom panel of Figure 2.6 shows that the low default intensities for bad risks in the years leading up to the financial crisis are attributed to the frailty component.

Finally, we treat contagion as an industry-level effect that gives rise to industry-specific default dynamics. However, contagion effects can also be present at the portfolio level. For example, a default of a financial firm can lead to the default of a firm in another industry. Our frailty factor will pick up these contagion effects across industries.

Figure 2.5: Frailty factor

We plot the estimated frailty risk factor from models M2 and M3. We report the conditional mean and conditional mode estimate. Graphed standard error bands refer to the conditional mean and are at a 0.95 confidence level.

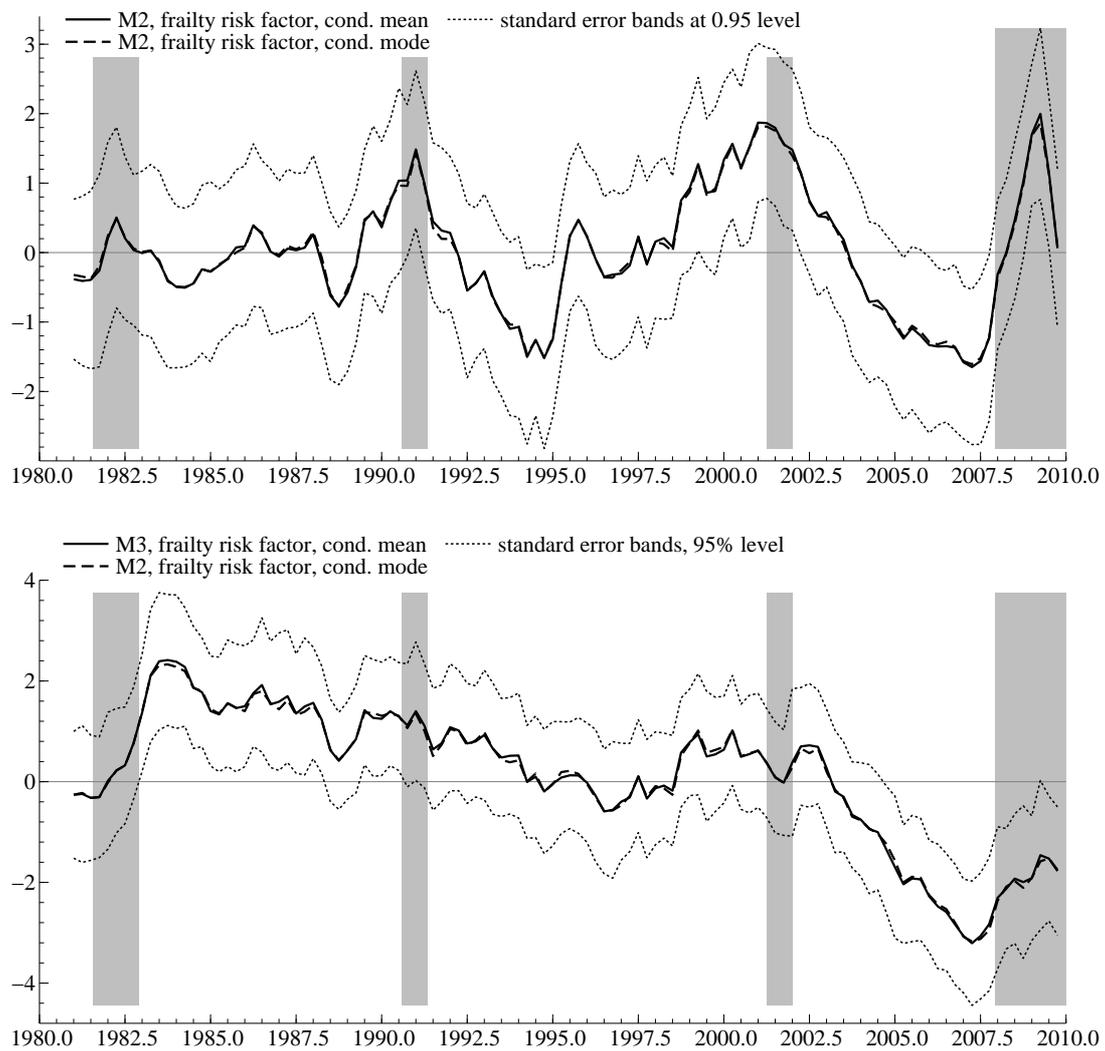
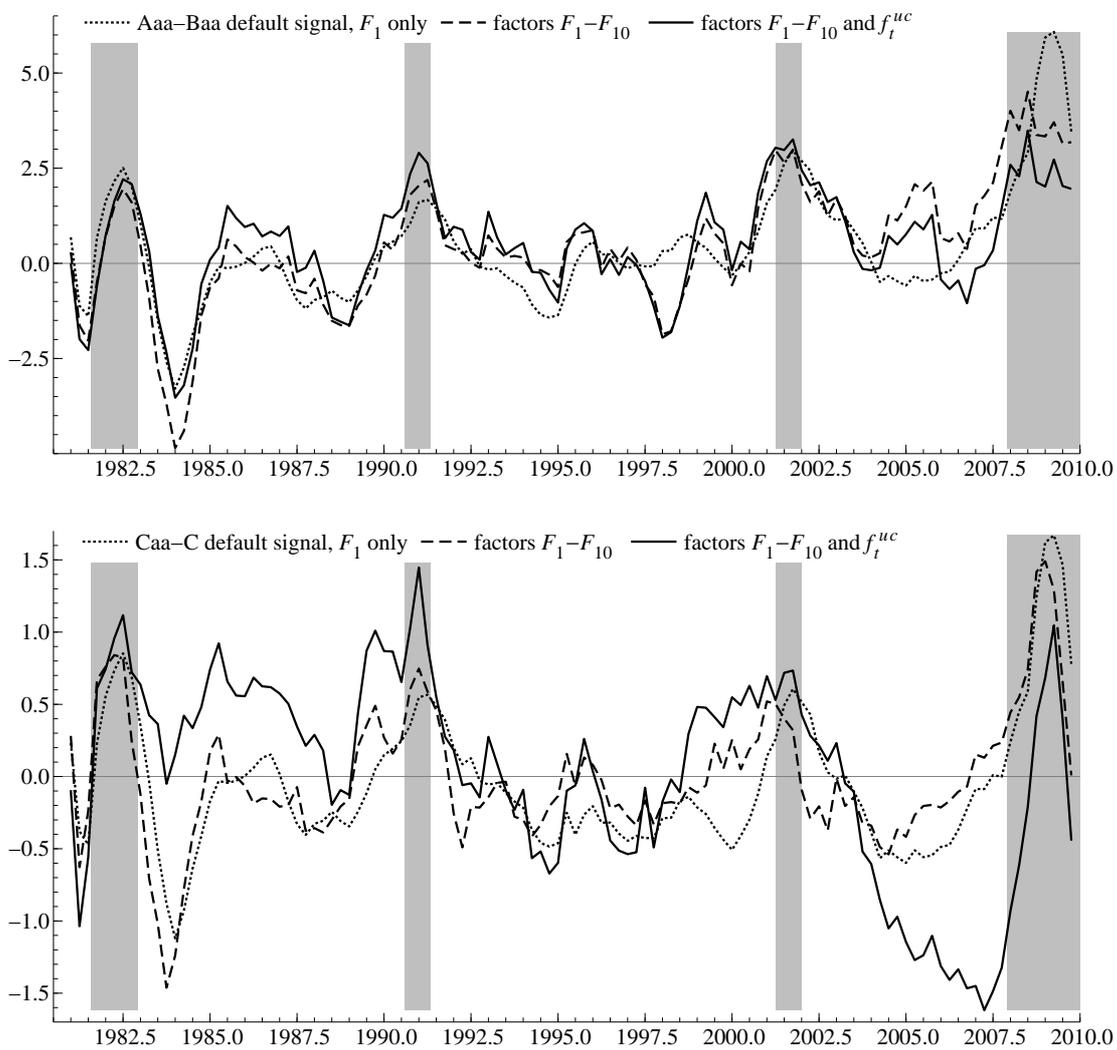


Figure 2.6: Smoothed default signals

The top and bottom figure plots smoothed default signals θ_{jt} for investment grade (Aaa-Baa) and low speculative grade (Caa-C) firms, respectively. The panels decompose the total default signal into estimated factors, scaled by their respective factor loadings (standard deviations). We plot variation due to the first principal component $\hat{F}_{1,t}$, all principal components $\hat{F}_{1,t}$ to $\hat{F}_{10,t}$, and all factors including the latent component \hat{f}_t^{uc} .



2.5.6 Out of sample forecasting accuracy

We compare the out-of-sample forecasting performance between models by considering a number of competing model specifications. Accurate forecasts are valuable in conditional credit risk management, for short-term loan pricing, and for credit portfolio stress testing. Also, out-of-sample forecasting is a stringent diagnostic check for modeling and analyzing time series. We present a truly out-of-sample forecasting study by estimating the parameters of the model using data upto a certain year and by computing the forecasts of the cross-sectional default probabilities for the next year. In this way we have computed our forecasts for the nine years of 2001, . . . , 2009.

The measurement of forecasting accuracy of time-varying intensities is not straightforward. Observed default fractions are only a crude measure of default conditions. We can illustrate this inaccuracy by considering a group of, say, 5 firms. Even if the default probability for this group is forecasted perfectly, it is unlikely to coincide with the observed default fraction of either 0, 1/5, 2/5, etc. The forecast error may therefore be large but it does not necessarily indicate a bad forecast. The observed default fractions are only useful when a sufficiently large number of firms are pooled in a single group. For this reason we pool default and exposure counts over age cohorts, and focus on two broad rating groups, i.e., (i) all rated firms in a certain industry, and (ii) firms in that industry with ratings *Ba* and below (speculative grade). The mean absolute error (MAE) and the root mean squared error statistic (RMSE) are computed as

$$\text{MAE}(t) = \frac{1}{D} \sum_{d=1}^D |\hat{\pi}_{d,t+4|t}^{an} - \bar{\pi}_{d,t+4}^{an}|, \quad \text{RMSE}(t) = \left(\frac{1}{D} \sum_{d=1}^D [\hat{\pi}_{d,t+4|t}^{an} - \bar{\pi}_{d,t+4}^{an}]^2 \right)^{\frac{1}{2}},$$

where index $d = 1, \dots, D$ refers to industry groups.

The estimated and realized annual probabilities are given by

$$\hat{\pi}_{d,t+4|t}^{an} = 1 - \prod_{h=1}^4 (1 - \hat{\pi}_{d,t+h|t}), \quad \bar{\pi}_{d,t+4}^{an} = 1 - \prod_{h=1}^4 \left(1 - \frac{y_{d,t+h}}{k_{d,t+h}} \right),$$

respectively, where $\hat{\pi}_{d,t+h|t}$, for $h = 1, \dots, 4$, are the forecasted quarterly probabilities for time $t + h$. To obtain the required default signals, we first forecast all factors $\hat{F}_t, \hat{f}_t^{uc}$ jointly using a low order vector autoregression and using the mode estimates of \hat{F}_t and \hat{f}_t^{uc} , in-sample. Although mode estimates of f_t^{uc} are indicated by \bar{f}_t^{uc} , in our forecasting study we integrate them in a Gaussian vector autoregression for which mode and mean estimates are the same. This vector autoregressive model takes into account that the factors F_t and f_t^{uc} are conditionally correlated with each other. Given the forecasts of

\hat{F}_t and \hat{f}_t^{uc} , we compute $\hat{\pi}_{d,t+h|t}$ using equations (2.2) and (2.3) and based on parameter estimates and mode estimates of the signal θ_{jt} .

Table 2.3 reports the forecast error statistics for five competing models. Model 0 does not contain common factors. It thus corresponds to the common practice of estimating default probabilities using long-term historical averages. We use a model with only baseline hazards and three well-used macros (industrial production growth, changes in the unemployment rate, and the credit spread (Aaa – Baa)) as our benchmark. The benchmark model is denoted as $M0(X_t)$.

Table 2.3: Out-of-sample forecasting accuracy

The table reports forecast error statistics associated with one-year ahead out-of-sample forecasts of time-varying point-in-time default probabilities/hazard rates. Error statistics are relative to a benchmark model $M0(X_t)$ with observed risk factors only, where X_t contains changes in industrial production, changes in unemployment rate, and the yield spread between Baa and Aaa rated bonds, see Section 2.5.6. We report mean absolute error (MAE) and root mean square error (RMSE) statistics for all firms (All) and speculative grade (SpG), respectively, based on all industry-group forecasts for the years 2001 – 2009. The relative MAEs are also given for all industry-group forecasts, for each year. Model M0 contains constant only. Models M1, M2, and M3 contain in addition the factors F_t , f_t^{uc} , and both F_t , f_t^{uc} , respectively. The models may also contain covariates as indicated.

Model		TOTAL	Ch.MAE	2001	2002	2003	2004	2005	2006	2007	2008	2009
M0: no factors	MAE	All	0.0%	1.05	0.66	1.62	1.08	1.01	1.04	1.02	1.06	0.80
	RMSE	SpG	-1.4%	1.01	0.65	1.58	1.08	1.01	1.04	1.03	1.05	0.76
		All	1.06	0.70	1.49	1.09	1.01	1.05	1.04	1.07	0.85	
		SpG	1.04	0.71	1.43	1.09	1.01	1.05	1.04	1.06	0.73	
M0: X_t, C_t	MAE	All	-0.8%	0.96	1.07	1.09	0.95	0.96	1.04	0.96	0.95	1.00
	RMSE	SpG	-1.3%	0.96	1.04	1.06	0.94	0.95	1.03	0.95	0.96	1.02
		All	0.97	1.12	1.06	0.96	0.98	1.05	0.96	0.97	1.07	
		SpG	1.00	0.97	1.07	1.03	0.95	0.97	1.04	0.95	0.97	1.07
M1: F_t, C_t	MAE	All	-9.4%	0.93	0.84	1.10	0.74	0.66	0.83	0.96	0.89	1.27
	RMSE	SpG	-10.2%	0.99	0.82	1.04	0.75	0.64	0.81	0.96	0.91	1.25
		All	0.89	0.77	1.04	0.77	0.71	0.85	0.97	0.88	1.34	
		SpG	0.96	0.77	1.00	0.78	0.69	0.85	0.97	0.89	1.31	
M2: f_t^{uc}, C_t	MAE	All	-38.7%	1.13	0.93	0.88	0.50	0.36	0.26	0.30	1.03	0.77
	RMSE	SpG	-38.2%	1.07	0.79	0.89	0.52	0.35	0.30	0.34	1.05	0.76
		All	1.15	0.88	0.80	0.55	0.41	0.28	0.35	1.17	0.76	
		SpG	1.08	0.82	0.82	0.57	0.40	0.33	0.38	1.17	0.72	
M3: $F_t, f_t^{uc}, \text{no } C_t$	MAE	All	-36.9%	0.90	0.58	0.62	0.45	0.37	0.39	0.32	1.19	1.18
	RMSE	SpG	-37.4%	0.92	0.57	0.74	0.44	0.37	0.42	0.31	1.28	1.08
		All	0.91	0.67	0.70	0.46	0.40	0.39	0.37	1.24	1.16	
		SpG	0.95	0.70	0.77	0.46	0.40	0.43	0.37	1.32	1.09	
M3: F_t, f_t^{uc}, C_t	MAE	All	-43.0%	0.95	0.58	0.61	0.37	0.35	0.31	0.29	1.10	0.98
	RMSE	SpG	-43.2%	0.94	0.56	0.73	0.36	0.35	0.35	0.29	1.18	0.86
		All	0.96	0.66	0.62	0.37	0.37	0.32	0.35	1.21	0.99	
		SpG	0.97	0.68	0.70	0.38	0.39	0.37	0.35	1.26	0.90	

Another version of Model 0 includes three observed variables instead of the common macro factors to forecast conditional default probabilities; they are changes in industrial production, changes in unemployment rate, and the yield spread between Baa and Aaa rated bonds. We label the benchmark model $M0(X_t)$. This approach is more common in the literature and here it serves as a more realistic benchmark. The results reported in Table 2.3 are based on out-of-sample forecasts from Models 1, 2, and 3, with their parameters replaced by their corresponding estimates as reported in Table 2.2.

As the main finding, ‘observed’ risk factors \hat{F}_t , the latent component f_t^{uc} , as well as the industry-specific contagion risk factors in C_t , each contribute to out-of-sample forecasting performance for default hazard rates, to different extents. Feasible reductions in forecasting error are substantial, and by far exceed the reductions achieved by using a few observed covariates directly.

The observed reduction in mean absolute forecasting error due to the inclusion of the three observed covariates from Model 0 is less than 2%. Using other observed risk factors provides similar results. Reductions in forecasting error increase when the observed covariates are replaced by principal components and are as high as 10% on average over the years 2001-2009. This finding shows that principal components from a large macro and finance panel can capture default dynamics more successfully.

Forecasts improve further when an unobserved component is added to the the principal components and contagion factors. Mean absolute forecasting errors then reduce to 43%. Reductions in MAE are most pronounced when frailty effects are highest. This is the case in 2002, when default rates remain high while the economy is recovering from recession, and years 2005-2007, when default conditions are substantially better than expected from macro and financial data. Reductions of more than 40% on average are substantial and have clear practical implications for the computation of capital requirements. It is also clear that the simple AR(1) dynamics for the frailty factor are too simplistic to capture the abrupt changes in common credit conditions during the crisis of 2008. As the frailty factor is negative over 2007, the forecast of default risk over 2008 based on the AR(1) dynamics is too low. In 2009, we find that the full model including frailty again does better than its competitors. To further improve the forecasting performance of the full model in crisis situations, one could extend the dynamic behavior of the frailty factor further to include non-linearity. This is left for future research.

2.6 Conclusion

We propose a novel non-Gaussian panel data time series model with regression effects to estimate and measure the dynamics of corporate default hazard rates. The model combines a non-Gaussian panel data specification with the principal components of a large number of macroeconomic covariates. The model integrates three different types of factors: common factors from macroeconomic and financial time series, an unobserved latent component for discrete default data, and ‘observed’ contagion factors at the industry level. At the same time we can include standard measures such as equity returns, volatilities, and ratings, in the model.

In an empirical application, the combined factors to capture a statistically significant share of the dynamics in the time series of disaggregated default counts. We find a large and significant role for a dynamic frailty component, even after accounting for more than 80% of the variation in more than 100 macroeconomic and financial covariates, and after controlling for contagion effects at the industry level. A latent component or frailty factor is thus needed to prevent a downward bias in the estimation of extreme default losses on portfolios of U.S. corporate debt. Our result also indicates that the presence of a latent factor may not be due to a few omitted macroeconomic covariates, but rather appears to capture different omitted effects at different times.

In an out-of-sample forecasting experiment, we obtain substantial reductions between 10% and 43% on average in mean absolute error when forecasting conditional point-in-time default probabilities using our factor structure. The forecasts from our model are particularly accurate in times when frailty effects are important and when aggregate default conditions deviate from financial and business cycle conditions. A frailty component implies additional default rate volatility, and may contribute to default clustering during periods of stress. Practitioners who rely on observed macroeconomic and firm-specific data alone may underestimate their economic capital requirements and crisis default probabilities as a result.

Chapter 3

Mixed measurement dynamic factor models

3.1 Introduction

We develop a novel latent dynamic factor model for panels of mixed measurement time series data. In this framework, observations may come from different families of parametric distributions, may be observed at different frequencies, and are dependent in the cross-section due to shared exposure to latent dynamic factors. Consider available data

$$y_t = (y_{1t}, \dots, y_{Nt})', \quad t = 1, \dots, T, \quad (3.1)$$

where each row $y_i = (y_{i1}, \dots, y_{iT})$, $i = 1, \dots, N$, comes from a different density. We are not only thinking of simple differences in means or variances. Instead, some time series may be discrete, whereas others are continuous. Some time series may be Gaussian, while others are non-negative durations, or count data obtained from point processes. Time series data from the exponential family is often of particular interest. This family includes many well-known distributions, such as the binomial, Poisson, Gaussian, inverse Gaussian, Gamma, and Weibull distribution. The results of this paper allow us to analyze the joint variation in mixed data from the above densities, in a latent dynamic factor model setting.

In the absence of non-Gaussian or mixed data, latent factors underlying a panel of time series data can be analyzed using either (i) the method of principal components in an approximate dynamic factor model framework, see e.g. Connor and Korajczyk (1986, 1988, 1993), Stock and Watson (2002, 2005), Bai (2003), and Bai and Ng (2002, 2007), (ii) estimation procedures based on frequency domain methods, see e.g. Sargent and Sims (1977), Geweke (1977), Forni, Hallin, Lippi, and Reichlin (2000, 2005), or (iii) filtering and

smoothing techniques in a state space framework, see e.g. Doz, Giannone, and Reichlin (2006), and Jungbacker and Koopman (2008). If data (3.1) come from different families of densities, however, none of the above methods can be used for parameter and factor estimation without modification. To the best of our knowledge, this paper is the first to present a likelihood-based analysis of a dynamic factor model for mixed measurement time series data. We refer to the model as the mixed-measurement dynamic factor model (MM-DFM).

In this paper, the main challenge is that the likelihood of the MM-DFM does not exist in closed form. Obviously, this hinders parameter and factor estimation and inference in a likelihood-based setting. We present three solutions to this problem. First, Shephard and Pitt (1997), Durbin and Koopman (2000), and Jungbacker and Koopman (2007) show that maximum likelihood inference and latent factor estimation can be achieved by Monte Carlo maximum likelihood methods based on importance sampling techniques. We cast the MM-DFM in state space form, and demonstrate that this approach can be extended to the MM-DFM setting. Recent applications of importance sampling in a non-Gaussian framework include Koopman, Lucas, and Monteiro (2008), Koopman and Lucas (2008), and Koopman, Lucas and Schwaab (2008, 2010).

Second, we consider a less complex observation driven model as an alternative to the parameter driven model in state space form. Here, the scaled score of the (local) log-likelihood function serves as the driving mechanism for the latent factors. This essentially eliminates the factor's second source of error. Creal, Koopman, and Lucas (2008) refer to such models as generalized autoregressive score (GAS) models. Effectively, this paper extends the class of GAS models to include a panel data model for observations from different families of parametric distributions (MM-GAS). Importantly, the likelihood exists in closed form, and can be maximized straightforwardly.

Third, we demonstrate that parameter and factor estimation in the MM-DFM framework can be performed by Bayesian techniques. Bayesian inference is particularly attractive when there is some prior information about the parameters. Also, Markov Chain Monte Carlo (MCMC) methods still work in settings with a very high dimensional factor space, where we may want to sample the factors in blocks.

As an example of a mixed-measurement and mixed-frequency panel data setting, we model the systematic variation in cross sections of corporate default counts, recovery rates on loans and bonds after default, and macroeconomic data. While defaults are discrete, the recovered percentages on the principal are continuous and bounded on the unit interval. Recovery values tend to be low precisely when defaults are high in an economic downturn, indicating important systematic covariation across different types

of data from different families of parametric distributions. In addition, recovery rates, default counts, and macroeconomic indicators are available at different frequencies. It is not difficult to think of further applications from e.g. the actuarial sciences or financial market microstructure research. Mixed-measurement models are useful whenever different families of statistical distributions are appropriate for different types of data, while they may be driven by related dynamics.

The remainder of this paper is as follows. We introduce the baseline mixed measurement dynamic factor model (MM-DFM) in Section 3.2, along with results regarding parameter estimation and signal extraction in this framework. We demonstrate how to speed up likelihood evaluations by collapsing observations, and address missing values. Section 3.3 introduces an observation driven MM-GAS alternative. Bayesian inference for the MM-DFM is treated in Section 3.4. Section 3.5 considers the estimation and forecasting of intertwined credit and recovery risk conditions. Section 3.6 concludes.

3.2 Mixed-measurement dynamic factor model

This section introduces a parameter driven latent dynamic factor model for variables from a broad range of densities, which we refer to as the mixed-measurement dynamic factor model (MM-DFM). Variables may be observed at different frequencies, such as monthly, quarterly, annually, etc.

3.2.1 Model specification

The mixed measurement dynamic factor model is based on a set of m dynamic latent factors that are assumed to be generated from a dynamic Gaussian process. For example, we can collect the factors into the $m \times 1$ vector f_t and assume a stationary vector autoregressive process for the factors,

$$f_{t+1} = \mu_f + \Phi f_t + \eta_t, \quad \eta_t \sim N(0, \Sigma_\eta), \quad t = 1, 2, \dots, \quad (3.2)$$

with the initial condition $f_1 \sim N(\mu, \Sigma_f)$. The $m \times 1$ mean vector μ_f , the $m \times m$ coefficient matrix Φ and the $m \times m$ variance matrix Σ_η are assumed fixed and unknown with the m roots of the equation $|I - \Phi z| = 0$ outside the unit circle and Σ_η positive definite. The $m \times 1$ disturbance vectors η_t are serially uncorrelated. The process for f_t is initialized by $f_1 \sim N((I - \Phi)^{-1} \mu_f, \Sigma_f)$ where $m \times m$ variance matrix Σ_f is a function of Φ and Σ_η or, more specifically, Σ_f is the solution of $\Sigma_f = \Phi \Sigma_f \Phi' + \Sigma_\eta$.

Conditional on a factor path $\mathcal{F}_t = \{f_1, f_2, \dots, f_t\}$, the observation $y_{i,t}$ of the i th variable at time t is assumed to come from a certain density given by

$$y_{i,t} | \mathcal{F}_t \sim p_i(y_{i,t}; \mathcal{F}_t, \psi), \quad i = 1, \dots, N. \quad (3.3)$$

For example, the observation $y_{i,t}$ could come from the exponential family of densities,

$$p_i(y_{i,t}; \mathcal{F}_t, \psi) = \exp\{a_i(\psi)^{-1} [y_{i,t} \theta_{i,t} - b_{i,t}(\theta_{i,t}; \psi)] + c_{i,t}(y_{i,t})\}, \quad (3.4)$$

with the signal defined by

$$\theta_{i,t} = \alpha_i + \sum_{j=0}^p \lambda'_{i,j} f_{t-j}, \quad (3.5)$$

where α_i is an unknown constant and $\lambda_{i,j}$ is the $m \times 1$ loading vector with unknown coefficients for $j = 0, 1, \dots, p$. The so-called link function in (3.4) $b_{i,t}(\theta_{i,t}; \psi)$ is assumed to be twice differentiable while $c_{i,t}(y_{i,t})$ is a function of the data only. The parameter vector ψ contains all unknown coefficients in the model specification including those in Φ , α_i and $\lambda_{i,j}$ for $i = 1, \dots, N$ and $j = 0, 1, \dots, p$. Scaling by a dispersion parameter $a_i(\psi)$ in (3.4) is not necessary for binary, binomial, Poisson, exponential, negative binomial, multinomial, and standard normal observations, as $a_i(\psi) = 1$ in these cases. Allowing for $a_i(\psi) \neq 1$ permits modeling observations from e.g. the Gamma, Gaussian, inverse Gaussian, and Weibull densities. In general, the results of Section 3.2.2 extend to densities $p_i(y_{i,t}; \mathcal{F}_t, \psi)$ which are twice differentiable with respect to their signal $\theta_{i,t}$, and $\partial^2 p_i(\cdot; \cdot) / \partial \theta_{i,t}^2 < 0$ to ensure positive implied variances. To enable the identification of all entries in ψ , we assume standardized factors in (3.2) which we enforce by the restrictions $\mu_f = 0$ and $\Sigma_f = I$ implying that $\Sigma_\eta = I - \Phi\Phi'$.

Conditional on \mathcal{F}_t , the observations at time t are independent of each other. It implies that the density of the $N \times 1$ observation vector $y_t = (y_{1,t}, \dots, y_{N,t})'$ is given by

$$p(y_t | \mathcal{F}_t, \psi) = \prod_{i=1}^N p_i(y_{i,t} | \mathcal{F}_t, \psi).$$

The MM-DFM model is defined by the equations (3.2), (3.3) and (3.5).

3.2.2 Estimation via Monte Carlo maximum likelihood

An analytical expression for the maximum likelihood (ML) estimate of parameter vector ψ for the MM-DFM is not available. Let $y = (y'_1, \dots, y'_T)'$ and $f = (f'_1, \dots, f'_T)'$ denote the vector of all the observations and factors, respectively. Let $p(y|f; \psi)$ be the density

of y conditional on f and let $p(f; \psi)$ be the density of f . The log-likelihood function is only available in the form of an integral

$$p(y; \psi) = \int p(y, f; \psi) \, df = \int p(y|f; \psi)p(f; \psi) \, df, \quad (3.6)$$

where f is integrated out. A feasible approach to computing this integral is provided by importance sampling; see, e.g. Kloek and van Dijk (1978), Geweke (1989) and Durbin and Koopman (2001). Upon computing the integral, the maximum likelihood estimator of ψ is obtained by direct maximization of the likelihood function using Newton-Raphson methods.

Importance sampling proceeds by finding a proposal distribution $g(f|y; \psi)$, called the importance density, which closely approximates $p(f|y; \psi)$ but has heavier tails. Assume that the conditions underlying the application of importance sampling hold, in particular that $g(f|y; \psi)$ is sufficiently close to $p(f|y; \psi)$ and simulation from $g(f|y; \psi)$ is feasible. Then a Monte Carlo estimate of the likelihood $p(y; \psi)$ can be obtained as

$$\tilde{p}(y; \psi) = g(y; \psi) M^{-1} \sum_{k=1}^M \frac{p(y|f^{(k)}; \psi)}{g(y|f^{(k)}; \psi)}, \quad f^{(k)} \sim g(f|y; \psi), \quad (3.7)$$

where M is a large number of draws. Density $g(y; \psi)$ is the likelihood of an approximating model which is employed to obtain the samples $f^{(k)} \sim g(f|y; \psi)$, see below. A derivation of (3.7) is provided in the appendix A1.

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

$$\tilde{y}_t = \theta_t + \varepsilon_t, \quad \varepsilon_t \sim N(0, \tilde{H}_t), \quad (3.8)$$

where the transition equation for θ_t is the same as in the original model of interest. The pseudo-observations \tilde{y}_t and covariance matrices \tilde{H}_t are chosen in such a way that the distribution $g(f|y; \psi)$ implied by the approximating state space model is sufficiently close to the distribution $p(f|y; \psi)$ from the original non-Gaussian model. Shephard and Pitt (1997) and Durbin and Koopman (1997) argue that \tilde{y}_t and \tilde{H}_t can be uniquely chosen such that the mode and curvature at the mode of $g(f|y; \psi)$ match the mode and curvature of $p(f|y; \psi)$ for a given value of ψ . The following algorithm shows how an approximating model can be obtained in a MM-DFM setting.

Algorithm 1: A linear Gaussian approximating model for the mixed measurement DFM

can be obtained by iterating the following steps until convergence. To this purpose define

$$\dot{p}_{i,t} = \left. \frac{\partial \log p_i(y_{i,t}|\theta_{i,t})}{\partial \theta_{i,t}} \right|_{\theta_{i,t}=\tilde{\theta}_{i,t}} \quad \text{and} \quad \ddot{p}_{i,t} = \left. \frac{\partial^2 \log p(y_{i,t}|\theta_{i,t})}{\partial \theta_{i,t} \partial \theta'_{i,t}} \right|_{\theta_{i,t}=\tilde{\theta}_{i,t}},$$

with $\theta_t = (\theta_{1,t}, \dots, \theta_{N_t})$ and $\theta = (\theta_1, \dots, \theta_T)$.

1. Initialize a guess $\tilde{\theta}$ of the mode.
2. Given a current guess $\tilde{\theta}$, compute $\tilde{y}_{i,t} = \tilde{\theta}_{i,t} - \tilde{p}_{i,t}^{-1} \dot{p}_{i,t}$, and $\tilde{H}_{i,t} = -\tilde{p}_{i,t}^{-1}$ for $i = 1, \dots, N$. Let $\tilde{y}_t = (\tilde{y}_{1t}, \dots, \tilde{y}_{N_t})'$ and $\tilde{H}_t = \text{diag}(\tilde{H}_{1,t}, \dots, \tilde{H}_{N_t})$, for $t = 1, \dots, T$.
3. With \tilde{y}_t and \tilde{H}_t from Step 2, apply the Kalman filter and smoother to the state space model (3.8) to obtain the smoothed estimates $\hat{\theta}_t$ for $t = 1, \dots, T$. Set $\tilde{\theta} = \hat{\theta}$ as the next guess for the solution to the mode. Return to Step 2 until convergence.

A derivation of the updating equations is provided in the appendix A2. A possible metric of convergence is the sum of absolute percentage change between $\hat{\theta}$ and $\tilde{\theta}$. We briefly describe how to implement this procedure for several examples; these will be useful in the application below.

Illustration 1: As an example for deriving the updating equations of Algorithm 1, we consider a univariate time series y_t , $t = 1, \dots, T$, from a Binomial distribution with time-varying success probability π_t and time-varying number of trials k_t . The log-density $\log p(y_t|\pi_t) = y_t[\log \pi_t - \log(1 - \pi_t)] + k_t \log(1 - \pi_t) + \log \binom{k_t}{y_t}$ can be rewritten in terms of the canonical/natural parameter $\theta_t = \log[\pi_t/(1 - \pi_t)]$ as $\log p(y_t|\theta_t) = y_t \theta_t - k_t \log[1 + \exp(\theta_t)] + \log \binom{k_t}{y_t}$. The signal θ_t is assumed to exhibit factor structure (3.5), i.e., θ_t is an affine function of factors.

Differentiating the log-density with respect to its signal gives $\dot{p}_t = y_t - k_t e^{\theta_t}/(1 + e^{\theta_t})$, and $\ddot{p}_t = -k_t e^{\theta_t}/(1 + e^{\theta_t})^2$. Given a value of parameters ψ , Algorithm 1 now implies that we can match densities $p(f|y; \psi)$ and $g(f|y; \psi)$ by iterating on steps 2 and 3 with $\tilde{H}_t = k_t^{-1} e^{-\theta_t} (1 + e^{\theta_t})^2$ and $\tilde{y}_t = \tilde{\theta}_t + \tilde{H}_t \left[y_t - k_t \frac{e^{\theta_t}}{1 + e^{\theta_t}} \right]$. After convergence, draws can be taken from the approximating density $g(f|y; \psi)$ to evaluate the likelihood as indicated in (3.7).

Illustration 2: Gaussian observations do not need to be updated because the approximating model and original model coincide. Consider Gaussian time series observations y_t with time-varying mean θ_t and fixed variance σ^2 . The mean θ_t may vary due to exposure to latent dynamic factors f_t , see (3.5). Differentiating the Gaussian log-density with respect to its signal gives $\dot{p}_t = \sigma^{-2}[y_t - \tilde{\theta}_t]$, and $\ddot{p}_t = -\sigma^{-2}$. As a result, the updating takes

the form $\tilde{H}_t = \sigma^2$, and $\tilde{y}_t = \tilde{\theta}_t + [y_t - \tilde{\theta}_t] = y_t$. The fact that no updating is necessary in this relevant case is fortunate, since it speeds up calculation of the approximating model. We also note here that Gaussian observations cancel in the calculation of the importance sampling weights, since the actual and approximating densities coincide.

Illustration 3: As a final example, we consider observations $0 < y_t < 1$ from a Beta(a, b) distribution, with $a, b > 0$. The first parameter a is often interpreted to capture mainly the ‘location’ of the observation according to $E[y_t] = a/(a + b)$, while the second parameter b may determine the ‘scale’ according to $\text{Var}[y_t] = ab/(a + b)^2(a + b + 1)$. Following this interpretation, we give a factor structure to the first parameter, $a = \theta_t$, to capture observed covariation with other time series of interest. The Beta log-density is given by $\log p(y_t|\theta_t; b) = -\log B(\theta_t; b) + (\theta_t - 1) \log y_t + (b - 1) \log(1 - y_t)$, with $\theta_t > 0, b > 0$, where $B(\theta_t; b) = \Gamma(\theta_t)\Gamma(b)/\Gamma(\theta_t + b)$ ensures that the density integrates to one. This implies $\dot{p}_t = \bar{\phi}(\theta_t + b) - \bar{\phi}(\theta_t) + \log y_t$ and $\ddot{p}_t = \bar{\phi}'(\theta_t + b) - \bar{\phi}'(\theta_t) < 0$, where $\bar{\phi}(x) = \Gamma'(x)/\Gamma(x)$ denotes the digamma function. The updating steps are formulated as above. Considering time variation in the second parameter is also possible.

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 (3.8). For a set of M draws of $g(f|y; \psi)$, the evaluation of (3.7) 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 (3.3), density $g(y|f; \psi)$ is based on the Gaussian density for $y_{i,t} - \mu_{it} - \theta_{i,t} \sim N(0, \sigma_{i,t}^2)$ (3.8) and $g(y; \psi)$ can be computed by the Kalman filter applied to (3.8), see Schweppe (1965) and Harvey (1989).

3.2.3 Estimation of the factors

Once an ML estimator is available for ψ , the estimation of the location of f can be based on importance sampling. It can be shown that

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

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

$$\tilde{f} = \frac{\sum_{k=1}^M w_k \cdot f^{(k)}}{\sum_{k=1}^M w_k}, \quad (3.9)$$

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

$$s_t^2 = \left(\sum_{k=1}^M w_k \cdot (f_t^{(k)})^2 \right) / \left(\sum_{k=1}^M w_k \right) - \tilde{f}_t^2, \quad (3.10)$$

with \tilde{f}_t the t th elements of \tilde{f} . A derivation of (3.9) and (3.10) is provided in Appendix A1. The availability of conditional variance estimates allows us to construct estimated standard error bands around the conditional mean of the factors.

As an alternative estimator of the latent factors f_t , we may obtain the conditional mode as given by

$$\bar{f} = \operatorname{argmax} p(f|y; \psi). \quad (3.11)$$

The conditional mode indicates the most probable value of the factors given the observations. In practice, it is obtained automatically as a by-product when matching the modes of densities $p(f|y; \psi)$ and $g(f|y; \psi)$, see Algorithm 1. In practice, \bar{f} and \tilde{f} are usually very close, see also Section 3.5.

3.2.4 Collapsing observations

A recent result in Jungbacker and Koopman (2008) states that it is possible to collapse a $[N \times 1]$ vector of (Gaussian) observations y_t into a vector of transformed observations y_t^l of lower dimension $m < N$ without compromising the information required to estimate factors f_t via the Kalman Filter and Smoother. This subsection adapts their argument to a nonlinear mixed-measurement setting. We focus on collapsing the artificial Gaussian data \tilde{y}_t with associated covariance matrices \tilde{H}_t , see (3.8) and (3.2).

Consider a linear approximating model for transformed data $\tilde{y}_t^* = A_t \tilde{y}_t$, for a sequence of invertible matrices A_t , for $t = 1, \dots, T$. The transformed observations are given by

$$\tilde{y}_t^* = \begin{pmatrix} \tilde{y}_t^l \\ \tilde{y}_t^h \end{pmatrix}, \quad \text{with } \tilde{y}_t^l = A_t^l \tilde{y}_t \text{ and } \tilde{y}_t^h = A_t^h \tilde{y}_t,$$

where time-varying projection matrices are partitioned as $A_t = [A_t^l : A_t^h]'$. We require (i) matrices A_t to be of full rank to prevent the loss of information in each rotation, (ii) $A_t^h \tilde{H}_t A_t^l = 0$ to ensure that observations \tilde{y}_t^l and \tilde{y}_t^h are independent, and (iii) $A_t^h Z_t = 0$ to ensure that y_t^h does not depend on f . Several such matrices A_t^l that fulfill these conditions can be found. A convenient choice is presented below. Matrices A_t^h can be constructed from A_t^l , but are not necessary for computing smoothed signal and factor estimates.

Given matrices A_t , a convenient model for transformed observations \tilde{y}_t^* is of the form

$$\begin{aligned} \tilde{y}_t^l &= A_t^l \theta_t + e_t^l, \\ \tilde{y}_t^h &= e_t^h \end{aligned}, \quad \begin{pmatrix} e_t^l \\ e_t^h \end{pmatrix} \sim NIID \left(0, \begin{bmatrix} \tilde{H}_t^l & 0 \\ 0 & \tilde{H}_t^h \end{bmatrix} \right),$$

where $\tilde{H}_t^l = A_t^l \tilde{H}_t A_t^{l'}$, $\tilde{H}_t^h = A_t^h \tilde{H}_t A_t^{h'}$, $\theta_t = Z f_t$, and Z contains the factor loadings. Clearly, the $[N - m]$ dimensional vector \tilde{y}_t^h contains no information about f_t . We can speed up computations involving the KFS recursions as follows.

Algorithm 2: Consider (approximating) Gaussian data \tilde{y}_t with time-varying covariance matrices \tilde{H}_t , and $N > m$. To compute smoothed factors f_t and signals θ_t ,

1. construct, at each time $t = 1, \dots, T$, a matrix $A_t^l = C_t Z' \tilde{H}_t^{-1}$, with C_t such that $C_t' C_t = \left(Z' \tilde{H}_t^{-1} Z \right)^{-1}$ and C_t upper triangular. Collapse observations as $\tilde{y}_t^l = A_t^l \tilde{y}_t$.
2. apply the Kalman Filter and Smoother (KFS) to the $[m \times 1]$ low-dimensional vector \tilde{y}_t^l with time-varying factor loadings $C_t^{-1'}$ and $\tilde{H}_t^l = I_m$.

This approach gives the same factor and signal estimates as when the KFS recursions are applied to the $[N \times 1]$ dimensional system for \tilde{y}_t with factor loadings Z and covariances \tilde{H}_t .

A derivation is provided in Jungbacker and Koopman (2008, Illustration 4). Collapsing observations in the MM-DFM involves a tradeoff. On the one hand, less observations need to be passed through the KFS after collapsing observations. This leads to savings in computing time. On the other hand, collapsing observations requires the Choleski decomposition of a (small) $[m \times m]$ matrix at each time $t = 1, \dots, T$, which is not required in a linear Gaussian dynamic factor model. As a result, the reductions in computing time depend on N , T , and m . Savings increase with N , and decrease with m and T .

3.2.5 Missing values due to mixed frequencies and forecasting

This section addresses the treatment of missing values. Missings arise easily when data is available at different sampling frequencies. Missing values also arise in out-of-sample forecasting at the end of the sample. For mixed frequency data, we suggest arranging the data on a grid at the higher frequency. For example, variables at a monthly and quarterly frequency can be arranged on a monthly grid. The quarterly series will then contain missing values. The precise arrangement may depend on whether data is a stock (point in time) or flow (a quantity over time, or average) measurement.

Missing values are accommodated easily in a state space approach. Most implementations of the Kalman filter (KF) and associated smoother (KFS) automatically assign a zero Kalman gain, zero prediction error, and large (infinite) prediction error variance to missing observations, see e.g. the implementation by Koopman, Shephard, and Doornik (2008). As a result, little extra effort is required. Some care must be taken when computing the importance sample weights $w_k = p(y|f^{(k)}; \psi)/g(y|f^{(k)}; \psi)$, $f^{(k)} \sim g(f|y; \psi)$. While $y = (y'_1, \dots, y'_T)'$ may contain many missing values, the (mode) estimates of the corresponding signals $\theta = (\theta'_1, \dots, \theta'_T)'$ and factors $f = (f'_1, \dots, f'_T)'$ are available for all data. Some bookkeeping is therefore required to evaluate $p(y|f; \psi)$ and $g(\tilde{y}|f; \psi)$ at the corresponding values of f , or θ .

Forecasting in the MM-DFM framework has several advantages over the two-step approach of e.g. Stock and Watson (2002b). First, forecasting factors and observations in the MM-DFM framework does not require the formulation of an auxiliary model. Parameter estimation, signal extraction, and forecasting occurs in a single step. In a two-step approach, factors are extracted from a large panel of predictor variables first, and a second step relates the variable of interest to the estimated factors. A simultaneous modeling approach (i) is conceptually straightforward, (ii) retains valid inference which is usually lost in a two step approach, and (iii) ensures implicitly that the extracted common factors are related to the variable of interest.

Forecasting factors is straightforward. Forecasts f_{T+h} , for $h = 1, 2, \dots, H$, can be obtained by treating future observations y_{T+1}, \dots, y_{T+H} as missing, and applying the estimation and signal extraction techniques of Section 3.2.2 to data (y_0, \dots, y_{T+H}) . The obtained conditional mean \tilde{f} and mode forecasts \bar{f} of the factors provide a location and maximum-probability forecast given observations, respectively. The mean (or median, mode) predictions of observations $(y_{T+1}, \dots, y_{T+H})$ can be obtained as nonlinear functions of $(f_{T+1}, \dots, f_{T+H})$.

3.3 Mixed measurement generalized autoregressive score models

This section introduces an observation driven alternative to the parameter driven MM-DFM by adjusting the factor (state) equation. We refer to Creal, Koopman, and Lucas (2008) who recently proposed a framework for observation-driven time-varying parameters models, which is referred to as generalized autoregressive score (GAS) models. This subsection extends the GAS family of models to include a dynamic factor model for

mixed measurement panel data (MM-GAS).

3.3.1 Model specification MM-GAS

The observation and signal equation of the MM-DFM and MM-GAS model coincide, i.e.,

$$y_{i,t} \sim p_i(y_{i,t}|\theta_{i,t}; \psi) \quad \theta_{i,t} = \alpha_i + \sum_{j=1}^p \lambda'_{i,j} f_{t-j}. \quad (3.12)$$

The observation densities are functions of a latent $m \times 1$ vector of factors that are assumed to come from a vector autoregressive specification. Instead of having their own source of error, the factors f_t in a GAS model are driven by the scaled score of the (local) log-density of y_t according to

$$f_{t+1} = \mu_f + \sum_{i=1}^{\bar{p}} A_i s_{t-i+1} + \sum_{j=1}^q B_j f_{t-j+1}, \quad (3.13)$$

where μ_f is a vector of constants, and coefficient matrices A_i and B_j are of appropriate dimension $[m \times m]$ for $i = 0, \dots, \bar{p} - 1$ and $j = 1, \dots, q$. The scaled score s_t is a function of past observations, factors, and unknown parameters. Unknown coefficients from A_i, B_j, μ_f , etc., are collected in a vector ψ . The scaled score is given by

$$s_t = S_t \nabla_t, \quad (3.14)$$

where

$$\nabla_t' = \frac{\partial \log p(y_t|\theta_t; \psi)}{\partial \theta_t} \frac{\partial \theta_t}{\partial f_t'}, \quad \text{and } S_t = E_{t-1}[\nabla_t \nabla_t']^{-1} = \mathcal{I}_t^{-1}, \quad (3.15)$$

such that the scaling matrix S_t is equal to the conditional Fisher information matrix. In most models of interest, the information matrix equality holds such that $S_t^{-1} = E_{t-1}[\nabla_t \nabla_t'] = -E_{t-1} \left[\frac{\partial^2 \log p(y_t|\theta_t; \psi)}{\partial f_t \partial f_t'} \right]$. The updating mechanism (3.14) for f_t is a Gauss-Newton iteration for each new observation y_t that becomes available. The updating equation is based on the (local) likelihood score and associated information matrix and therefore exploits the full density structure to update the factors. Given that factors are common across observations from different families of densities, scaling by (3.15) gives an automatic and model consistent way to weight the information provided by different observations.

3.3.2 Maximum likelihood estimation

Parameter and factor estimation by maximum likelihood for the MM-GAS model is simpler and less computationally demanding compared to the Monte Carlo methods required

in the state space framework. The likelihood can be built recursively since current factors f_t , while stochastic, are perfectly predictable given past values of observations, factors, and coefficients ψ . Unknown parameters can be estimated by maximizing the log-likelihood

$$\max l(\psi) = \sum_{t=1}^T l(\psi; y^t, \mathcal{F}_t), \quad (3.16)$$

where $y^t = (y_1, \dots, y_t)$, $\mathcal{F}_t = (f_1, \dots, f_t)$, and $l(\psi; y^t, \mathcal{F}_t) = \log p(y_t | \mathcal{F}_t, \psi)$ for observed values y_t . Factors and likelihood increments are computed at each time t according to (3.13) and (3.16). Analytical derivatives for the score of the log-likelihood (3.16) can be obtained, but are usually complicated. In practice we therefore prefer to maximize the likelihood based on numerical derivatives. For a discussion whether standard asymptotic results apply, we refer to Creal, Koopman, and Lucas (2008, Section 3).

As in the MM-DFM setting of Section 3.2, we need to impose certain restrictions to ensure the identification of all parameters in ψ . As is common in factor models, a rotation of the factors by an invertible matrix, along with an inverse rotation of the factor loadings, yields an observationally equivalent model. As a result, we impose $\mu_f = 0$ in (3.13), and restrict certain factor loadings $\lambda_{i,j}$ in (3.12) to be rows of the corresponding identity matrix. We need to restrict as many rows of factor loadings as there are common factors in the model. Restricting the factor loadings identifies the unknown parameters in (3.13). This requirement is related to the scaling of $\Sigma_\eta = I - \Phi\Phi'$ in (3.2) to identify the factor loadings in the parameter driven framework.

We can still estimate (filtered) factors in the MM-GAS framework when portions of the panel are missing. For an unbalanced panel, we need to distinguish which part of the data is observed at each time $t = 1, \dots, T$. The increment in the log-likelihood for y_t , the score vector ∇_t , and scaling matrix S_t take contributions only from observed data. As in the state space model, forecasts f_{T+h} for $h = 1, 2, \dots, H$ can be obtained by treating future observations y_{T+1}, \dots, y_{T+H} as missing. Alternatively, factors may also be forecast as a random walk based on the latest filtered values (which implies $A = I_m$).

3.4 Bayesian inference

Bayesian inference is an alternative approach to overcome the complication that the likelihood of the MM-DFM is not available in closed form. Parameter and factor estimation by Markov Chain Monte Carlo (MCMC) is most useful when researchers have prior information about parameters. In addition, MCMC may still work in the (rare) cases in which the importance sampler does not appear to possess a variance. In that case we would like

to sample the factors in smaller chunks. An MCMC loop for parameters and factors can be constructed as follows.

3.4.1 Sampling the latent factors

We sample latent factors from its conditional density, i.e., $f^{(i)} \sim p(f|y, \psi^{(i-1)}, f^{(i-1)})$. For instance, this can be achieved by the simulation smoothing algorithm of Durbin and Koopman (2002) after constructing a Gaussian approximating model as in Section 2.2. The simulation smoother runs the Kalman Filter forward, the Kalman smoothing algorithm backward, and another run forward to simulate all factors in one step. The simulated factors come from a Gaussian proposal density, and are accepted with a probability that is related to the importance sampling weight for this draw.

In case sampling all factors at once appears too ambitious, a single site (or block) random walk Metropolis sampler can be used. A new proposal value for factors $f_t^{(i)}$ can be constructed from previously sampled values $f_t^{(i-1)}$ by adding a vector of error terms. Adding Gaussian errors yields a (symmetric) Gaussian proposal density. The proposed new value is accepted with probability $\alpha_t = \min\left(\frac{P(f_t^{(i)})}{P(f_t^{(i-1)})}, 1\right)$, $t = 2, \dots, T - 1$, where $\frac{P(f_t^{(i)})}{P(f_t^{(i-1)})}$ is the likelihood ratio of the proposed sample and the previous sample. This likelihood ratio depends on data y and neighboring previous samples, $f_{t-1}^{(i-1)}$ and $f_{t+1}^{(i-1)}$. The boundaries $t = 1$ and $t = T$ can be handled similarly. If rejected, $f_t^{(i+1)} = f_t^{(i)}$. The scales of the random walk Metropolis sampler are tuned to achieve a roughly 35% acceptance rate.

3.4.2 Sampling factor loadings and autoregressive parameters

We sample the parameters from their conditional density, i.e. $\psi^{(i)} \sim p(\psi|y, f^{(i)}, \psi^{(i-1)})$. Under the assumption that the factor loadings $\lambda_{i,j}$ and signal intercepts $\lambda_{0,i}$ in (3.5) have a conjugate normal prior, they can be obtained by regression. They are sampled from a normal distribution in a Gibbs sampling step.

The factor autoregressive parameters are restricted to lie in the unit interval. A beta prior for these parameters is standard, but not conjugate. Samples of the autoregressive parameters can be obtained in a random-walk Metropolis step.

3.5 Intertwined credit and recovery risk

Evidence from many countries in recent years suggests that collateral values and recovery rates on corporate defaults are volatile and, moreover, that they tend to go down just when the number of defaults goes up in economic recessions, see Altman, Brady, Resti, and Sironi (2003) for a survey. The inverse relationship between recovery rates and default rates has traditionally been neglected by credit risk models, treating the recovery rate as either constant or as a stochastic variable independent from the probability of default. It is now widely recognized that a failure to take these dependencies into account leads to incorrect forecasts of the loss distribution and the derived capital allocation, see Schuerman (2006).

According to the current Basel proposal, banks can opt to provide their own recovery rate forecasts for the calculation of regulatory capital, see Basel Committee on Banking Supervision (2004). As a result there is an immediate need for statistical modeling, in particular for the supervisory agencies who need to evaluate the banks' models. In credit risk practice, default counts are frequently modeled as conditionally binomial random variables, where default probabilities depend on unobserved systematic risk factors, see McNeil, Frey and Embrechts (2005, Chapter 9) and McNeil and Wendin (2007). Recovery rates take values on the unit interval, and may be given either a beta-distribution, as in CreditMetrics (2007) and Gupton and Stein (2005), or a logit-normal distribution, as in Düllmann and Trapp (2004) and Rösch and Scheule (2005).

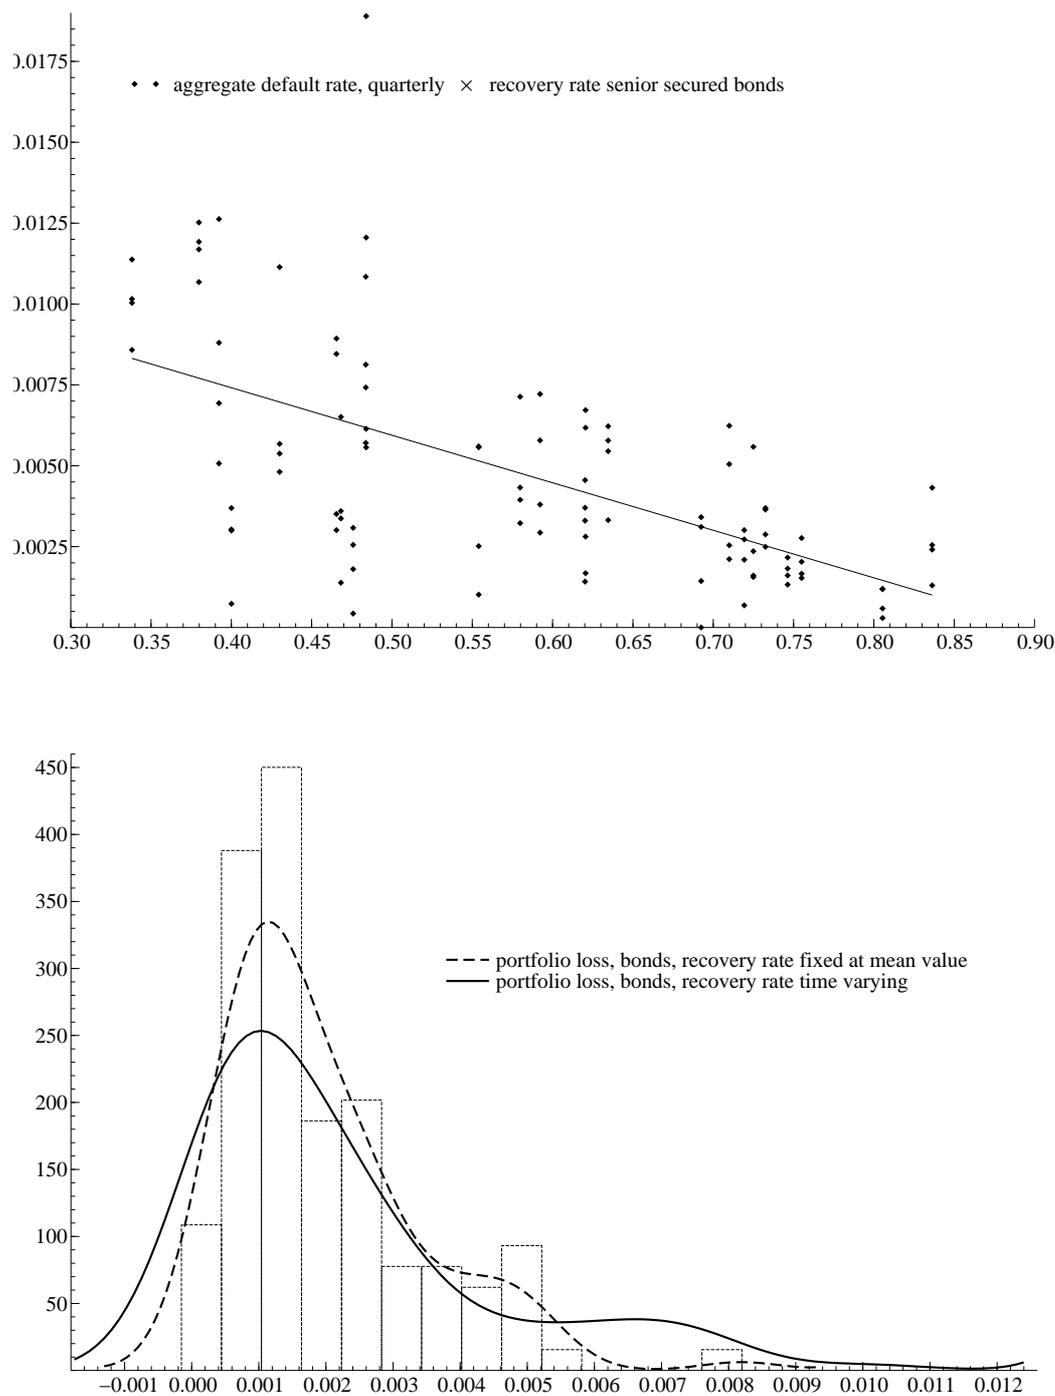
The top graph in Figure 3.1 illustrates the inverse relationship between observed default risk conditions and recovery rates. In bad times (high default rates), recoveries tend to be low, and vice versa. The effect of systematic recovery risk on the credit loss portfolio is illustrated in the bottom graph in Figure 3.1. The unconditional loss distribution refers to a setting where loans are given to each firm in the Moody's database at the beginning of each quarter from 1982Q1 to 2008Q4. The Figure shows a histogram of the portfolio losses due to corporate defaults (i) when the recovery rate is held constant at its mean value, and (ii) when the recovery rates vary inversely with defaults, as observed in the data. Systematic recovery risk implies that credit losses become more extreme: good times become better (thicker left tail), and worse times become worse (thicker right tail). Clearly, neglecting recovery risk leads to an underestimation of risk.

3.5.1 Data and mixed measurement model equations

Figure 3.2 contains time series plots of, from top to bottom, quarterly default counts of investment grade rated firms, quarterly default counts for firms with a speculative grade

Figure 3.1: Portfolio Loss Distributions with and without systematic recovery risk

The scatterplot in the top panel plots observed quarterly default rates for Moody's rated firms against average senior secured bond recovery rates over time. The regression line indicates an inverse relationship. The bottom panel presents a histogram of scaled historical default rates (the unconditional portfolio loss distribution) with and without systematic recovery rate risk. The panel compares the unconditional loss density (i) when recovery rates are held fixed at their mean value, and (ii) when historical recoveries vary inversely with the default rates.



rating, annual recovery rates for collateralized bank loans, recovery rates for senior secured bonds¹, changes in the US unemployment rate, and the negative of the US industrial production growth rate. The macroeconomic indicators are standardized to zero mean and unit variance. The observations are denoted $d_{j,t}$, $r_{j,t}$, and $x_{j,t}$, respectively, where $j = 1, 2$. Macroeconomic data from January 1982 to December 2008 is obtained from the Fred St. Louis online database. Rating and default data is from Moody's. Figure 3.2 exhibits the clear inverse relationship between defaults and recovery rates.

No recovery rates are reported for senior secured bonds in 1984 and 1993, due to a lack of informative default events in that year. Loan recovery rates are available only from 1990 onwards, yielding a time series of 19 annual observations. Fortunately, missing values are easily accommodated using the results in Section 3.2.5.

A parsimonious model for mixed measurement data $y_t = (d'_t, r'_t, x'_t)'$, with common exposure to latent autocorrelated risk factors f_t , is given by

$$d_{j,t}|f_t \sim \text{Binomial}(k_{j,t}, [1 + e^{-\theta_{j,t}}]^{-1}), \quad r_{j,t}|f_t \sim \text{Beta}(a_{j,t}, b_j), \quad x_{j,t}|f_t \sim N(\mu_{j,t}, \sigma_j^2), \quad (3.17)$$

where $[1 + e^{-\theta_{j,t}}]^{-1} = \pi_{j,t}$ denotes a time-varying default probability within the unit interval. Location parameters for each observation are given by

$$\theta_{j,t} = c_{\theta,j} + \beta'_j f_t, \quad a_{j,t} = c_{a,j} + \gamma'_j f_t, \quad \mu_{j,t} = c_{\mu,j} + \delta'_j f_t,$$

where $c_{\theta,j}$, $c_{a,j}$, $c_{\mu,j}$ are intercept terms and β_j , γ_j , δ_j are factor loadings. Unknown coefficients and factors can be estimated as outlined in Section 3.2.

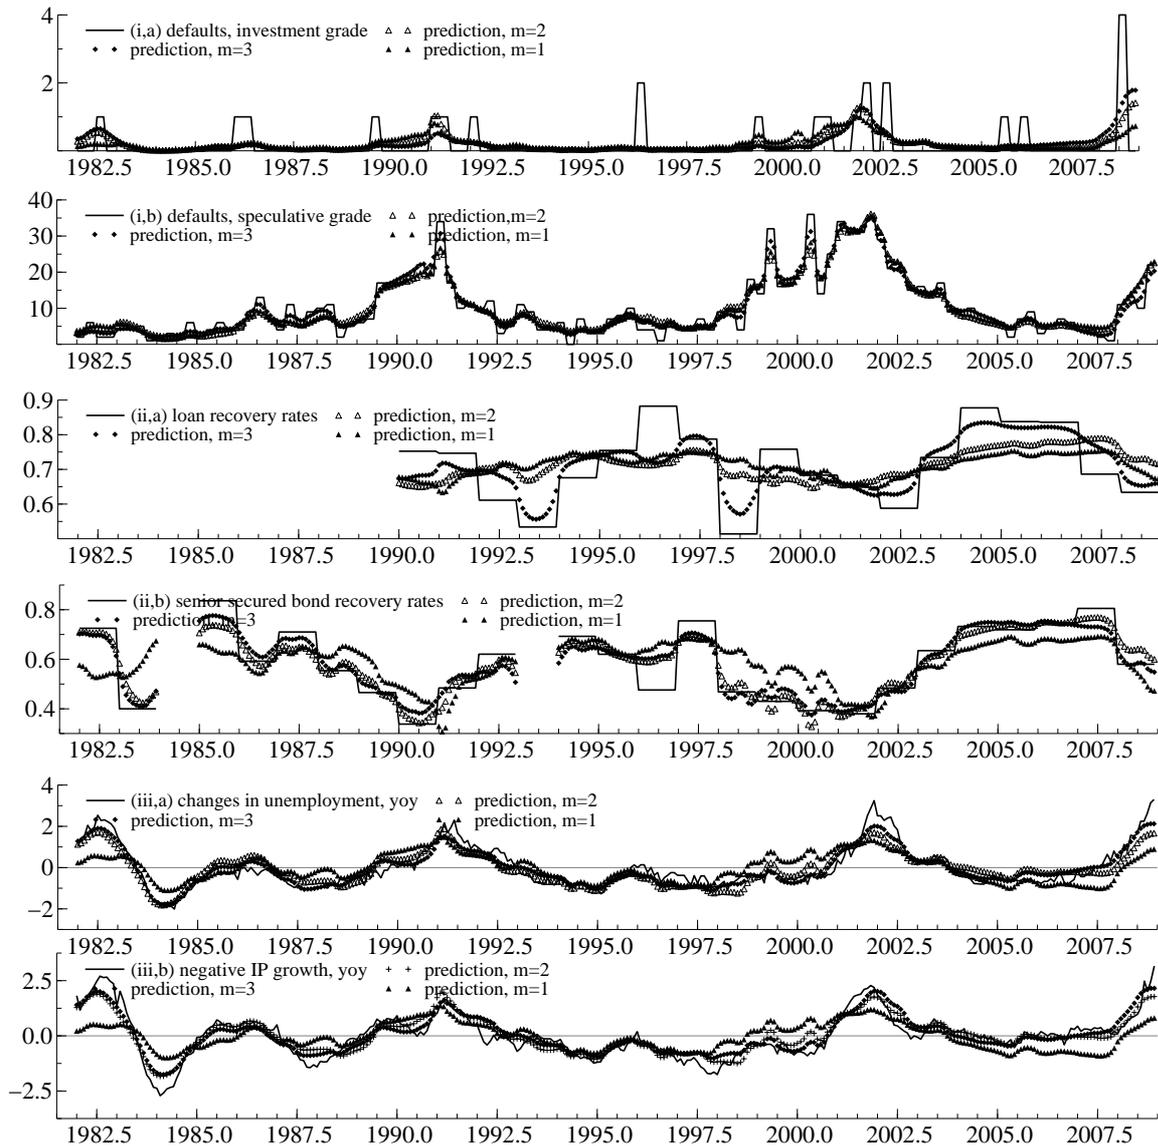
3.5.2 Major empirical findings

Figure 3.2 compares in-sample predictions for defaults, bond and loan recovery rates, and business cycle data to observed data. The single factor MM-DFM ($m = 1$) already gives an acceptable fit to the default counts and bond recovery rates. However, the fit to loan recovery rates and macroeconomic data is less satisfactory. This discrepancy may indicate that systematic default and recovery rate risk is related to, but different from, standard business cycle risk. This would confirm the related findings in Das, Duffie, Kapadia, and Saita (2007) and Bruche and Gonzalez-Aguado (2009). Extending the dimensionality of f_t yields a better fit, in particular for the macroeconomic indicators and bond recovery

¹Bond recovery rates are defined as the ratio of the market value of the bonds to the unpaid principal, one month after default, averaged across the bonds that default in a given year.

Figure 3.2: MM-DFM: Actual vs predicted values

The figure plots the actual versus predicted values of (i) default counts of firms rated investment grade and speculative grade, respectively, (ii) bank loan recovery rates, and recovery rates for senior secured bonds, and (iii) changes in the unemployment rate, yoy, and negative changes in industrial production. Defaults are quarterly data, recovery rates are annual data, and macro data is monthly data. Predicted values are obtained from a model specification with $m = 2$ and $m = 3$ factors, respectively.



rates.

The corresponding plots for the MM-GAS model are reported in Figure 3.3. For the current data, the observation driven alternative is able to replicate the in-sample fit of the MM-DFM.

Figure 3.3: MM-GAS: Actual vs predicted values

The figure plots the actual versus predicted values of (i) default counts of firms rated investment grade and speculative grade, respectively, (ii) bank loan recovery rates, and recovery rates for senior secured bonds, and (iii) changes in the unemployment rate, yoy, and negative changes in industrial production. Defaults are quarterly data, recovery rates are annual data, and macro data is monthly data. Predicted values are from a multi-factor MM-GAS model specification with $m = 3$ factors.

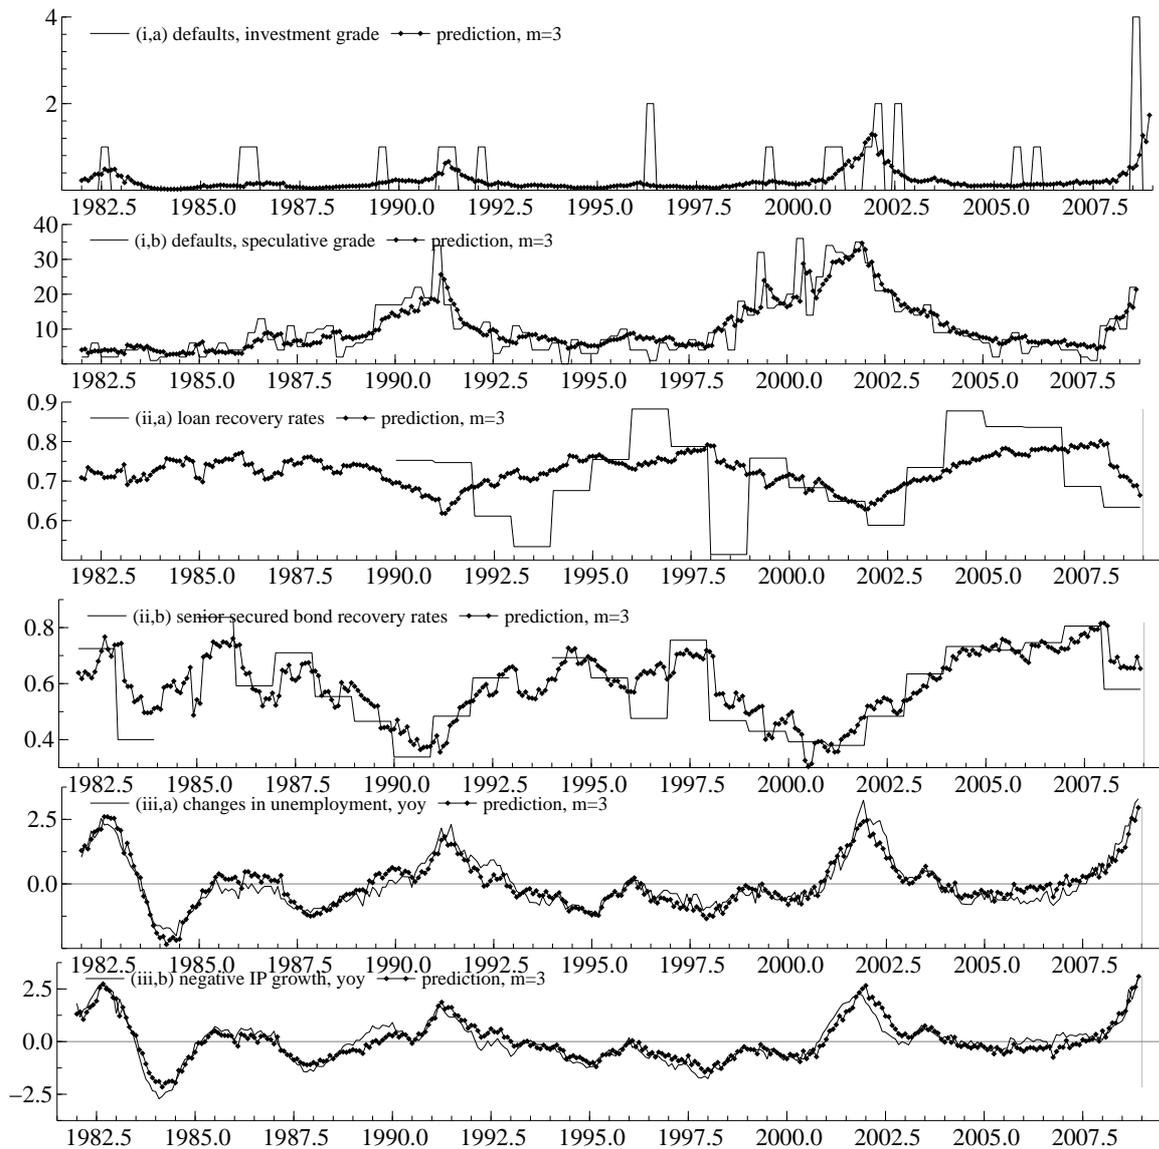
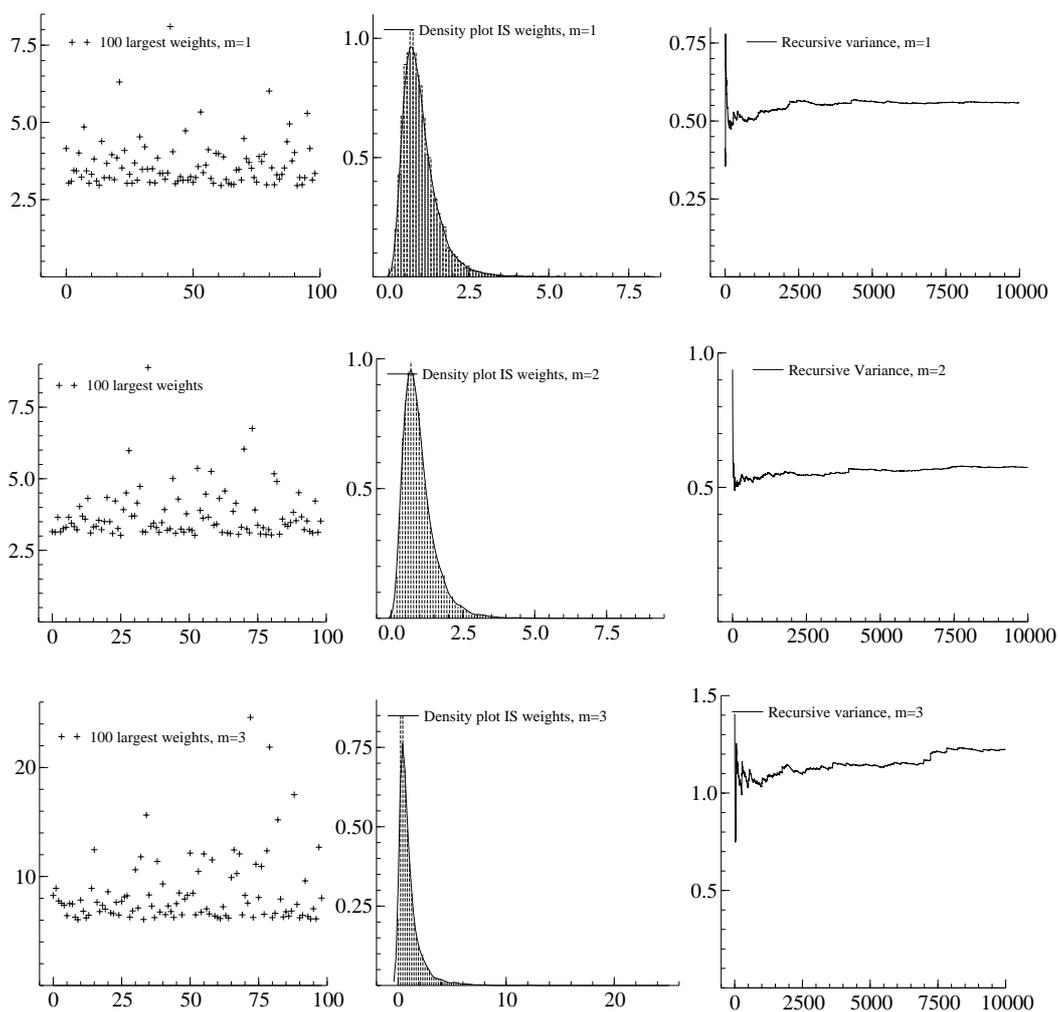


Figure 3.4: Importance sampling weights

The figure presents the largest 100 importance sampling weights, a density plot, and a recursive variance estimate for a total of 10000 simulated weights. The rows correspond to an empirical model specification with $m = 1$, $m = 2$, and $m = 3$ latent factors, respectively.



When estimating MM-DFM models, we assume that the assumptions underlying the application of importance sampling hold. In particular, $g(f|y; \psi)$ needs to approximate $p(f|y; \psi)$ sufficiently closely to ensure that the importance sampler possesses a variance. This guarantees a square root speed of convergence and asymptotic normality of the importance sampling estimators, see Geweke (1989). Some graphical diagnostics are presented in Figure 3.4. We present the largest 100 (log) importance sampling weights, a density plot, and a recursive variance estimate for 10000 importance sampling weights associated with models with $m = 1, 2, 3$ factors. There is no indication that a few extremely large weights dominate. The recursive variance estimates appear to converge. The largest weight accounts for less than 1% of the total sum of weights in all cases. However, the weights appear to become less well-behaved as more factors are added. Statistical tests for a finite variance are presented in Koopman, Shephard, and Creal (2009).

Figure 3.5 presents the conditional mean and conditional mode estimates for the three latent dynamic factors underlying the predictions in Figure 3.2. Both factor estimates are extremely close. The MM-GAS factors track the reported common factors from the MM-DFM.

3.5.3 Out of sample evaluation

This section compares the out-of-sample predictions of several models for mixed measurement data. We consider four model which differ widely in their degree of sophistication. We consider

1. a random walk forecast, assuming the last year's rates will remain the same,
2. a low order unrestricted vector autoregression, VAR (2), fitted on quarterly data for default rates, recovery rates, and macroeconomic time series. Missing data is replaced straightforwardly by its last known values.
3. the parameter driven MM-DFM, estimated by state space methods for different values of m . Recovery rates are fitted using time-varying parameter versions of the beta and logit-normal distribution.
4. several observation driven MM-GAS models, for different values of m .

Each model is used to produce an out-of-sample forecast of the (i) default rate for both investment grade and speculative grade rated issuers over the next year, (ii) loan and senior secured bond recovery rates for defaulted debt over the next year, and (iii) the annual change in US industrial production.

Figure 3.5: Latent factor estimates

The figure plots the estimates for three latent factors from a multi-factor ($m = 3$) MM-DFM and MM-GAS model specification, respectively. We report the conditional mean and mode estimates for the MM-DFM (left), and filtered factors for the MM-GAS (right). Standard error bands for the conditional mean of the factors are at a 0.95 confidence level.

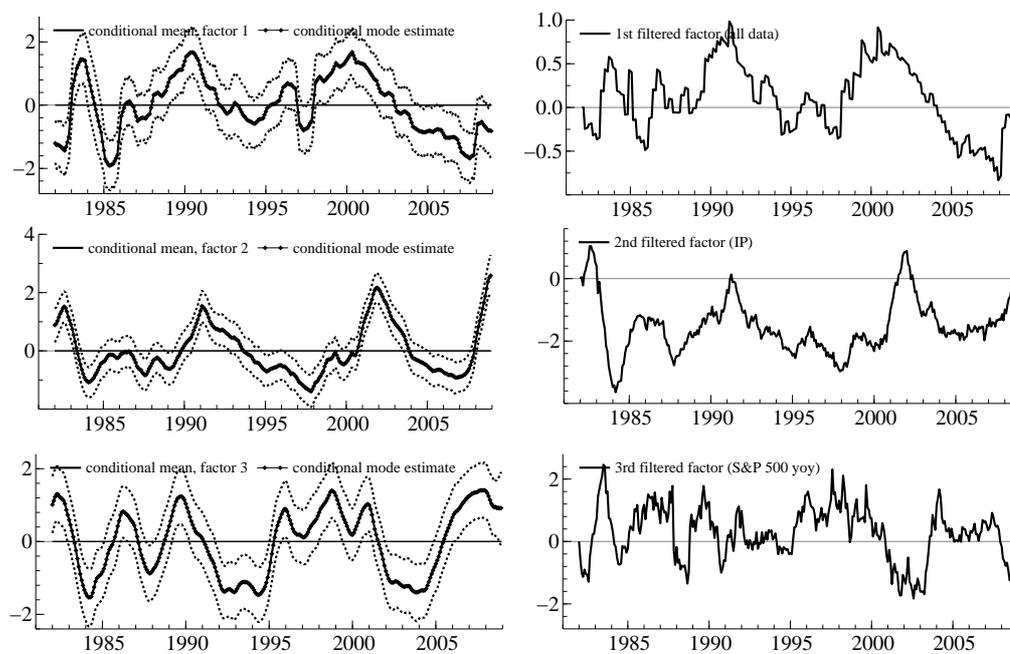


Table 3.1: Out of Sample Prediction Errors

The table presents the mean absolute error (MAE) and root mean square error (RMSE) statistics associated with out-of-sample point forecasts from different competing models.

MAE stats					
Model	IG def rate	SG def rate	loan rr	bond rr	ip growth
RW	1.09	0.49	0.11	0.14	2.03
VAR(2)	0.68	0.39	0.12	0.15	1.73
GAS(m=1)	0.88	0.43	0.13	0.16	1.77
GAS(m=2)	0.64	0.45	0.14	0.13	1.77
GAS(m=3)	0.75	0.41	0.15	0.14	1.94
GAS(m=4)	0.78	0.39	0.15	0.14	1.95
GAS(m=5)	0.76	0.32	0.13	0.17	1.94
SS(m=1, b)	0.69	0.49	0.12	0.18	1.72
SS(m=2, b)	0.77	0.46	0.13	0.16	2.07
SS(m=1, ln)	0.71	0.49	0.13	0.18	1.75
SS(m=2, ln)	0.74	0.44	0.14	0.14	1.98
SS(m=3, ln)	0.76	0.53	0.13	0.16	2.03
1/4 combi	0.78	0.42	0.12	0.15	1.80

RMSE stats					
Model	IG def rate	SG def rate	loan rr	bond rr	ip growth
RW	1.46	0.57	0.13	0.17	2.68
VAR(2)	0.85	0.48	0.14	0.18	2.31
GAS(m=1)	1.08	0.51	0.15	0.20	2.30
GAS(m=2)	0.81	0.53	0.15	0.18	2.30
GAS(m=3)	0.91	0.46	0.17	0.18	2.45
GAS(m=4)	0.92	0.43	0.17	0.18	2.44
GAS(m=5)	0.93	0.41	0.15	0.21	2.44
SS(m=1, b)	0.86	0.55	0.13	0.21	2.43
SS(m=2, b)	0.94	0.53	0.14	0.18	2.71
SS(m=1, ln)	0.86	0.55	0.14	0.20	2.45
SS(m=2, ln)	0.89	0.50	0.15	0.17	2.63
SS(m=3, ln)	0.93	0.61	0.16	0.18	2.63
1/4 combi	0.96	0.45	0.13	0.16	2.40

Table 3.1 presents the mean absolute error (MAE) and root mean square error (RMSE) statistics associated with one year ahead forecasts. Simple models, such as the VAR(2) and the Random Walk, do relatively well in forecasting. This holds in particular for random walk forecasts for the recovery rates, and the VAR forecasts of the default rate.

The MM-GAS model does at least as well as the more complex MM-DFM when predicting default rates. It also beats the Random Walk forecasts for default rates. This means that the increase in model tractability and estimation speed of the MM-GAS model compared to the MM-DFM does not come at the cost of reduced forecasting power. Extending the dimensionality of f_t for the factor models (MM-DFM and MM-GAS) tends to help for the prediction of some variables (speculative grade default rates, loan and bond recovery rates), but not for others (investment grade default rates, annual IP growth). Table 3.1 further suggests that the conditionally logit-normal and beta specifications for recovery rates do approximately equally well in prediction. The beta density seems slightly better for $m = 1$, while the logit-normal specification is better for $m = 2$. Both choices are comparable based on their out-of-sample performance.

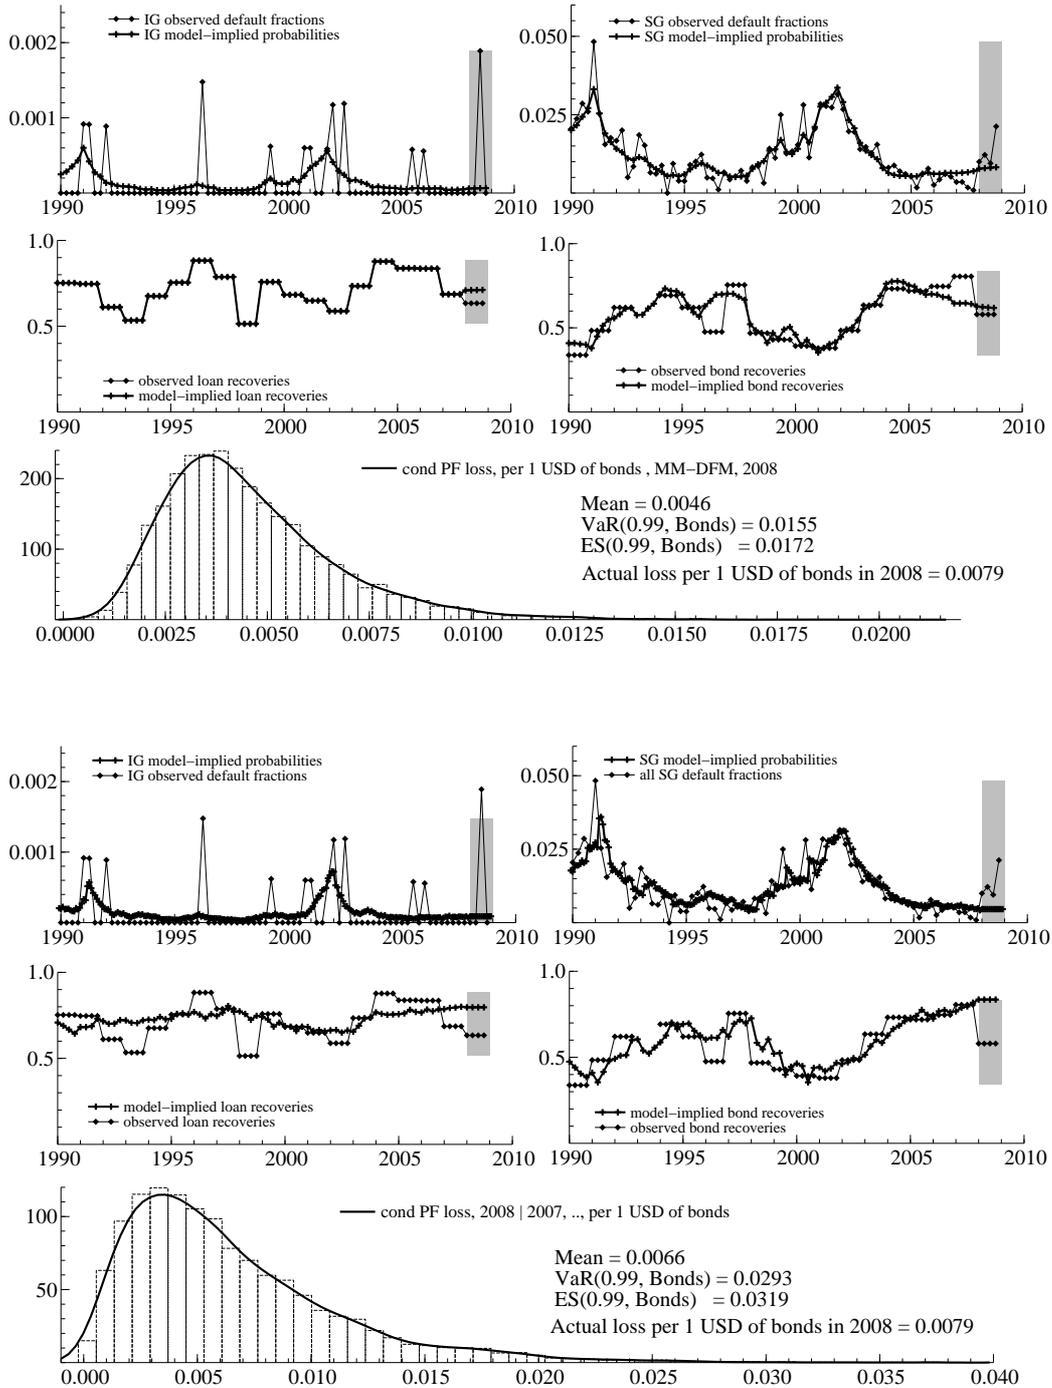
A combined forecast from four models [the Random Walk, the VAR(2), a MM-GAS model with five factors, and a MM-DFM with one factor, equal weighting] is often among the best three forecasts, and never among the three worst forecasts. The combined forecast has low prediction RMSEs, in particular for both recovery rates and speculative grade default rates. We conclude that the combined forecasts from two relatively simple (random walk, VAR(2)) and two sophisticated (MM-DFM, MM-GAS) models appears to give good joint forecasts of default rates and recovery rates.

Figure 3.6 plots the out-of-sample point forecasts for default and recovery rates for the recession year 2008. The forecasted levels are similar for the MM-DFM and MM-GAS model. In this particular case, the MM-DFM delivers better joint forecasts of defaults and recoveries than the GAS model. This finding can be explained by the fact that we have restricted (identified) two GAS factors to load mostly on macro data.

Figure 3.6 also presents the predictive density for the 2008 credit portfolio loss, conditional on macro, default, and recovery data from 1982 to 2007. The simulated predictive density from the MM-GAS model gives a wider confidence interval for the portfolio loss, and larger associated risk measures. It is therefore more conservative in this case. Both the MM-DFM and MM-GAS model imply capital buffers that would have ensured the solvency of a financial institution during that recession year.

Figure 3.6: Out-of-sample forecasts for 2008

The two panels plots the 1982 to 2007 in-sample predictions for quarterly investment and speculative grade default rates, and loan and bond recovery rates. Out-of-sample point forecasts for 2008 are based on a multi-factor ($m = 3$) MM-DFM (top panel) and MM-GAS model (bottom panel), respectively. We also plot the simulated out-of-sample predictive density for the portfolio credit loss based on bond default and recovery rate data.



3.6 Conclusion

We introduced a new latent dynamic factor model framework (MM-DFM) for time series observations from different families of parametric distributions and mixed sampling frequencies. As the main complication, the likelihood does not exist in closed form for this class of models. We therefore present simulation-based approaches to parameter and factor estimation in this framework. We also propose a less complex observation driven alternative to the parameter driven original model, for which the likelihood exists in closed form. Missing values arise due to mixed frequencies and forecasting, and can be accommodated straightforwardly in either the MM-DFM and MM-GAS framework. In an empirical application of the mixed-measurement framework we model the systematic variation in US corporate default counts and recovery rates from 1982 - 2008. We estimate and forecast intertwined default and recovery risk conditions, and demonstrate how to obtain the predictive credit portfolio loss distribution. While the MM-GAS model is simpler and computationally more efficient than the MM-DFM, we do not find that its reduced complexity comes at the cost of diminished out-of-sample prediction accuracy.

A1. Derivation of importance sampling estimators

Equations (3.7), (3.9) and (3.10) are derived below. Using importance sampling to estimate parameters and factors in nonlinear non-Gaussian models is not new, we refer to Shephard and Pitt (1997), and Durbin and Koopman (1997, 2000). For given parameters ψ , consider the estimation of the mean of an arbitrary function of the factors, $x = x(f)$, where $f = (f'_1, \dots, f'_T)'$, conditional on mixed measurement data $y = (y'_1, \dots, y'_T)'$,

$$\tilde{x} = \mathbb{E}[x(f)|y] = \int x(f)p(f|y; \psi) \, df.$$

There is no analytical solution for this problem. Denoting a suitable Gaussian importance density by $g(f|y; \psi)$,

$$\tilde{x} = \int x(f) \frac{p(f|y; \psi)}{g(f|y; \psi)} g(f|y; \psi) \, df = \mathbb{E}_g \left[x(f) \frac{p(f|y; \psi)}{g(f|y; \psi)} \right] = \frac{g(y; \psi)}{p(y; \psi)} \mathbb{E}_g \left[x(f) \frac{p(f, y; \psi)}{g(f, y; \psi)} \right], \quad (\text{A.18})$$

where \mathbb{E}_g denotes expectation with respect to $g(f|y; \psi)$. Setting $x(f) \equiv 1$ gives

$$1 = \frac{g(y; \psi)}{p(y; \psi)} \mathbb{E}_g \left[\frac{p(f, y; \psi)}{g(f, y; \psi)} \right], \quad (\text{A.19})$$

and thus

$$p(y; \psi) = g(y; \psi) \mathbb{E}_g \left[\frac{p(f, y; \psi)}{g(f, y; \psi)} \right], \quad (\text{A.20})$$

The Monte Carlo estimator (3.7) is the empirical counterpart to (A.20). It is of the same form as the estimator presented in Durbin and Koopman (1997). A law of large numbers, such as Khinchin's WLLN, ensures convergence under relatively weak conditions, see Geweke (1989).

Dividing (A.20) by (A.19) yields

$$\tilde{x} = \frac{\mathbb{E}_g [x(f)w(f, y; \psi)]}{\mathbb{E}_g [w(f, y; \psi)]}, \quad (\text{A.21})$$

where

$$w(f, y; \psi) = \frac{p(f, y; \psi)}{g(f, y; \psi)} = \frac{p(y|f; \psi) p(f; \psi)}{g(y|f; \psi) g(f; \psi)} = \frac{p(y|f; \psi)}{g(y|f; \psi)}.$$

The last equality uses the fact that the marginal distribution of the state is Gaussian, $p(f; \psi) = g(f; \psi)$. The weights $w_k = p(y|f^{(k)}; \psi)/g(y|f^{(k)}; \psi)$, $f^{(k)} \sim g(f|y; \psi)$ are i.i.d. by construction. Choices $x(f) = f$ and $x(f) = f^2$ in (A.21) give an expression for the first two conditional moments of f . A law of large numbers implies convergence of the empirical counterparts in (3.9) and (3.10). ■

A2. Derivation of Algorithm 1

We adapt a general argument for non-Gaussian models in state space form to the MM-DFM setting, compare Durbin and Koopman (2001), p. 192. For original work on importance sampling in a non-Gaussian framework we refer to Shephard and Pitt (1997), and Durbin and Koopman (1997, 2000). The dependence of observation densities on unknown parameters ψ is suppressed. The linear Gaussian approximating model is of the form (3.8) and (3.2). Let $g(f|y)$ and $g(f, y)$ be generated by the Gaussian approximating model, and let $p(f|y)$ and $p(f, y)$ be the corresponding densities as generated by the

mixed model (3.2), (3.3) and (3.5). We seek artificial data \tilde{y}_t and variances \tilde{H}_t such that the densities $g(f|y)$ and $p(f|y)$ have the same mode \bar{f} . The initialization condition for the unobserved factors is given by their stationary distribution, $g(f_1) = N(0, I_m)$. The (non-diffuse) initialization of factors and the time-invariance of MM-DFM system matrices Φ, Σ_η, \dots , simplify the exposition.

In the Gaussian model, the joint density $g(f, y)$ is given by

$$\log g(f, y) = \text{const} - \log g(f_1) - \frac{1}{2} \sum_t^T (f_{t+1} - \Phi f_t)' \Sigma_\eta^{-1} (f_{t+1} - \Phi f_t) - \frac{1}{2} \sum_t^T (y_t - Z f_t)' H_t^{-1} (y_t - Z f_t),$$

where $y_t \sim N(\theta_t, H_t)$ and signals are expressed as $\theta_t = Z f_t$. The conditional mode of $\log g(f|y) = \log g(f, y) - \log g(y)$ can be obtained as the solution to the first order condition

$$\frac{\partial \log g(f, y)}{\partial f_t} = (d_t - 1) f_1 - d_t \Sigma_\eta^{-1} (f_t - \Phi f_{t-1}) + \Phi' \Sigma_\eta^{-1} (f_{t+1} - \Phi f_t) + Z' H_t^{-1} (y_t - Z f_t) = 0, \quad (\text{A.22})$$

where $t = 1, \dots, T$, $d_1 = 0$ and $d_t = 1$ for $t = 2, \dots, T$, together with $\Sigma_\eta^{-1} (f_{T+1} - \Phi f_T) = 0$. Since $g(f|y)$ is Gaussian, the conditional mode \bar{f} is equal to the conditional mean $\bar{f} = E[f|y]$. The conditional mean is calculated efficiently by the Kalman filter and smoother (KFS), see e.g. Durbin and Koopman (2001), Chapter 4. It follows that the KFS recursions solve equation (A.22).

Assuming that the MM-DFM is sufficiently well-behaved, the mode of $\log p(f|y) = \log p(f, y) - \log p(y)$ is the solution to the vector equation

$$\partial \log p(f, y) / \partial f = 0, \quad (\text{A.23})$$

where $\log p(f, y) = \text{const} + \sum_{t=1}^T \log p(\eta_t) + \sum_{t=1}^T \log p(y_t | \theta_t)$, and $\eta_t = f_{t+1} - \Phi f_t$ as above. Thus, condition (A.23) becomes

$$\frac{\partial \log p(f, y)}{\partial f} = (d_t - 1) f_1 + d_t \frac{\partial \log p(\eta_{t-1})}{\partial \eta_{t-1}} - \Phi' \frac{\partial \log p(\eta_t)}{\partial \eta_t} + Z' \frac{\partial \log p(y_t | \theta_t)}{\partial \theta_t} = 0, \quad (\text{A.24})$$

where $d_1 = 0$ and $d_t = 1$ for $t = 2, \dots, T$. The first three terms of (A.24) and (A.22) are identical. The difference in the last terms is due to the observation component in the joint densities. It remains to linearize the last term of (A.24). Recall that

$$\dot{p}_t = \left. \frac{\partial \log p(Y_t | \theta_t)}{\partial \theta_t} \right|_{\theta_t = \tilde{\theta}_t} \quad \text{and} \quad \ddot{p}_t = \left. \frac{\partial^2 \log p(Y_t | \theta_t)}{\partial \theta_t \partial \theta_t'} \right|_{\theta_t = \tilde{\theta}_t},$$

such that a first-order expansion about $\tilde{\theta}_t$ gives approximately

$$\partial \log p(y_t | \theta_t) / \partial \theta_t = \dot{p}_t + \ddot{p}_t (\theta_t - \tilde{\theta}_t). \quad (\text{A.25})$$

Substituting (A.25) in the last term of (A.24) gives the linearized form $Z' (\dot{p}_t + \ddot{p}_t \theta_t - \ddot{p}_t \tilde{\theta}_t)$. To obtain a form which coincides with the last term in (A.22), choose

$$\tilde{y}_t = \tilde{\theta}_t - \ddot{p}_t^{-1} \dot{p}_t \quad \text{and} \quad \tilde{H}_t = -\ddot{p}_t^{-1}. \quad (\text{A.26})$$

These are the required updating equations. All elements in $y = (y'_1, \dots, y'_T)'$ are independent after conditioning on the corresponding signal $\theta = (\theta'_1, \dots, \theta'_T)'$. This implies that \tilde{H}_t is diagonal for all

$t = 1, \dots, T$. As a result, each observation can be updated individually. ■

A3. MM-GAS equations for credit risk model

We discuss the formulation of the MM-GAS model for the empirical application considered in Section 3.5. We consider the case of mixed measurements $y_t = (d_t', r_t', x_t')'$, where d_t is binomial, r_t is logit-normal, and x_t is Gaussian with time-varying parameters. The observations are dependent in the cross section since parameters depend on common factors.

$$d_{j,t}|f_t \sim \text{Binomial}(k_{j,t}, [1 + e^{-\theta_{j,t}}]^{-1}), \quad r_{j,t}|f_t \sim \text{Logit-normal}(\bar{\mu}_{j,t}, \bar{\sigma}_j^2), \quad x_{j,t}|f_t \sim \text{Normal}(\mu_{j,t}, \sigma_j^2),$$

where j indexes the cross section. The time-varying parameters depend on common factors as

$$\theta_{j,t} = c_{\theta,j} + Z_{d,j}f_t, \quad \bar{\mu}_{j,t} = c_{\bar{\mu},j} + Z_{r,j}f_t, \quad \mu_{j,t} = c_{\mu,j} + Z_{x,j}f_t.$$

The log-density for the observed variables y_t combines the multivariate normal, the binomial, and the logit-normal density. If all data is observed at time t , the local log-density is given by

$$\begin{aligned} \log p(y_t|f_t, \psi) &= \text{const} + \sum_{j=1}^{n_1} d_{j,t}\theta_{j,t} - k_{j,t} \log[1 + \exp(\theta_{j,t})] \\ &\quad - 0.5 \sum_{j=1}^{n_2} 2 \log[r_{j,t}/(1 - r_{j,t})] + \log \bar{\sigma}_j^2 + \bar{\sigma}_j^{-2} (\log[r_{j,t}/(1 - r_{j,t})] - \bar{\mu}_{j,t})^2 \\ &\quad - 0.5 \sum_{j=1}^{n_3} \log \sigma_j^2 + \sigma_j^{-2} (x_{j,t} - \mu_{j,t})^2, \end{aligned}$$

where n_1 , n_2 , and n_3 are the dimensions of d_t , r_t , and x_t . As the log-density, the score and information matrix for the factors f_t also depend on which data is observed at time t .

$$\begin{aligned} \nabla_t &= \sum_{j=1}^{n_1} (d_{j,t} - k_{j,t}[1 + \exp(-\theta_{j,t})]^{-1}) Z'_{d,j} \\ &\quad + \sum_{j=1}^{n_2} \bar{\sigma}_j^{-2} (\log[r_{j,t}/(1 - r_{j,t})] - \bar{\mu}_{j,t}) Z'_{r,j} + \sum_{j=1}^{n_3} \sigma_j^{-2} (x_{j,t} - \mu_{j,t}) Z'_{x,j}. \end{aligned}$$

The (inverse of) the scaling matrix $S_t^{-1} = E_{t-1}[\nabla_t \nabla_t']$ is given as

$$S_t^{-1} = Z'_d \Sigma_{d,t} Z_d + Z'_r \bar{\Sigma}^{-1} Z_r + Z'_x \Sigma^{-1} Z_x,$$

where $\Sigma_{d,t} = \text{diag}(\pi_{1,t}(1 - \pi_{1,t})k_{1,t}, \dots, \pi_{n_1,t}(1 - \pi_{n_1,t})k_{n_1,t})$, $\bar{\Sigma} = \text{diag}(\bar{\sigma}_1^2, \dots, \bar{\sigma}_{n_2}^2)$, and $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_{n_3}^2)$. In case data is missing at time t , the respective contributions to the sums are zero.

Chapter 4

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

4.1 Introduction

In this paper we test three competing explanations for systematic variation in default rates using a new methodological framework. 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 US 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 macro economic 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.

A first line of literature attributes the additional variation in default intensities to

an unobserved dynamic component, also known as a frailty factor. The discussion of frailty factors in the credit risk literature is fairly recent, see Das et al. (2007), McNeil and Wendin (2007), Koopman, Lucas, and Monteiro (2008), Koopman and Lucas (2008), Koopman, Lucas, and Schwaab (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, such as firms' expectations about future business conditions and the trust in the accuracy of public accounting information, see Duffie et al. (2009).

A second line of literature puts forward contagion as a relevant factor for additional default clustering. It refers to the phenomenon that a defaulting firm weakens other firms with which it has business links, see the discussion in Giesecke (2004) and Giesecke and Azizpour (2008). Contagion effects may dominate potentially offsetting competitive effects at the intra-industry level, see e.g. Lang and Stulz (1992). More detailed work by Jorion and Zhang (2007b) suggests that credit contagion may depend on the type of bankruptcy. A Chapter 11 bankruptcy is contagious ("bad"), while a Chapter 7 bankruptcy is competitive ("good"). Lando and Nielsen (2008) screen hundreds of default histories in the Moody's database 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.

It is not known to what extent the three different explanations (macro, frailty, industry/contagion) 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 (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.

To address this question, we decompose the systematic variation in corporate defaults into its different constituents as suggested in the literature. For this purpose, we develop a new methodological framework in which default rate volatility at the rating and industry level is attributed to macro, frailty, and industry effects simultaneously. The attractive feature of our framework is threefold. First, it allows us to combine standard continuous

time series (such as business cycle proxies, financial market conditions, and interest rates) with discrete series such as default counts. Second, and in contrast to earlier models, we can include a substantive number of macro controls to account for the different components of macro economic 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 forecasting, e.g. in the context of computing adequate capital requirements.

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, 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, in particular, 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 diagnostic 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) for the United States, and the European Systemic Risk Board (ESRB) for the European Union.

The remainder of this paper is organized as follows. Section 4.2 introduces our general methodological framework. Section 4.3 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 4.4. Section 4.5 concludes.

4.2 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

$$y_t = (y_{1t}, \dots, y_{Jt}, y_{J+1,t}, \dots, y_{J+N,t})', \quad (4.1)$$

for $t = 1, \dots, T$, 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}, \dots, 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 $j = J+1, \dots, J+N$, 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 $f'_t = (f_t^{m'}, f_t^{d'}, f_t^{i'})$. 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,

$$f_t = \Phi f_{t-1} + \eta_t, \quad t = 1, 2, \dots, \quad (4.2)$$

with the coefficient matrix Φ diagonal and with $\eta_t \sim N(0, \Sigma_\eta)$. More complex dynamics than (4.2) 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 $f_1 \sim N(0, \Sigma_0)$. 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 $j = 1, \dots, J$. 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 $j = J + 1, \dots, J + N$.

4.2.1 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

$$y_{jt}|\mathcal{F}_t \sim p_j(y_{jt}|\mathcal{F}_t; \psi), \quad (4.3)$$

where $\mathcal{F}_t = \{f_t, f_{t-1}, \dots\}$ and ψ is a vector of fixed and unknown parameters that include, for example, the elements of Φ and Σ_η in (4.2). 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

$$\theta_{jt} = \alpha_j + \lambda'_j f_t, \quad (4.4)$$

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 (4.4). In particular, we specify the signals by

$$\theta_{jt} = \lambda_{0j} + \beta'_j f_t^m + \gamma'_j f_t^d + \delta'_j f_t^i, \quad \text{for } j = 1, \dots, J, \quad (4.5)$$

$$\theta_{jt} = \lambda_{0j} + \beta'_j f_t^m, \quad \text{for } j = J + 1, \dots, J + N. \quad (4.6)$$

The signal specification in (4.6) 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 (4.5) 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 (4.5) and (4.6) 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 extent 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 (4.2) 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 (4.2). We aim to estimate ψ by maximum likelihood. For this purpose, we numerically maximize the likelihood function as given by

$$p(y; \psi) = \int p(y, f; \psi) df = \int p(y|f; \psi)p(f; \psi)df, \quad (4.7)$$

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 (4.7) is not known analytically, and we therefore rely on numerical methods. The likelihood function (4.7) 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.

4.2.2 Decomposition of non-Gaussian variation

Once the model is 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.

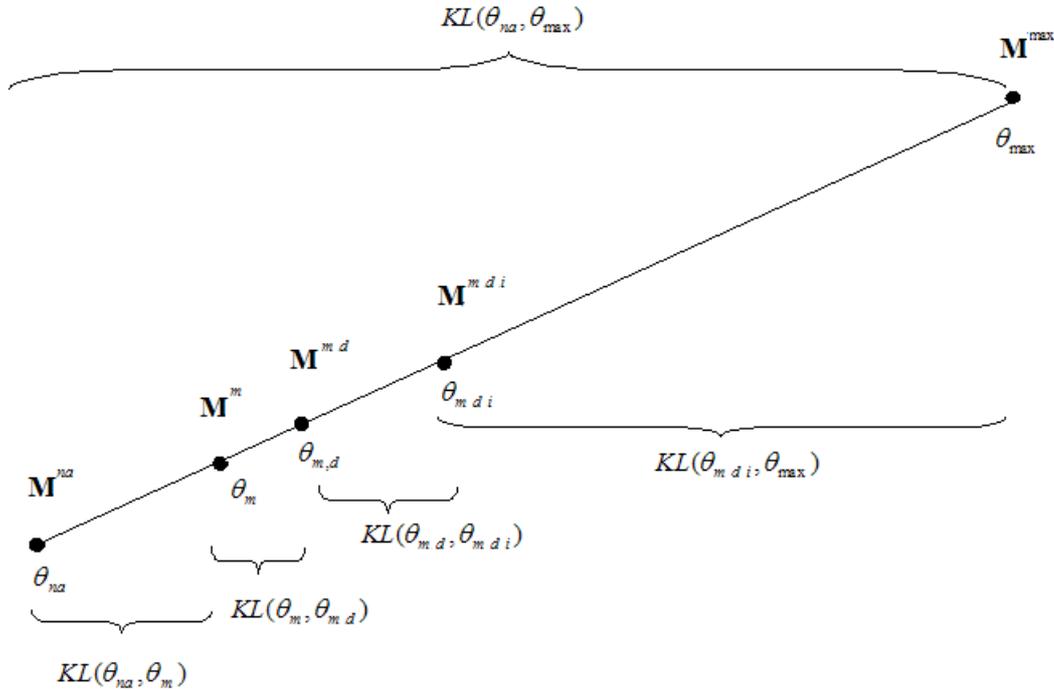
Our distance measure for the pseudo- R^2 is the Kullback-Leibler (KL) divergence, which is defined as

$$KL(\theta_1, \theta_2) = 2 \int [\log p_{\theta_1}(y) - \log p_{\theta_2}(y)] p_{\theta_1}(y) dy. \quad (4.8)$$

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

Figure 4.1: 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.



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 4.1 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 , $M^{m,d}$, and $M^{m,d,i}$. These models contain an increasing collection of latent factors. Model M^{na} does not contain any factors, while models M^m , $M^{m,d}$, and $M^{m,d,i}$ 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.

For each model specification M^{na} , M^m , $M^{m,d}$, $M^{m,d,i}$, and M^{max} , we obtain a time series of fitted log-odds from (4.5). The factor estimates \hat{f}_t and the parameter estimates

$\hat{\beta}_j$, $\hat{\gamma}_j$ and $\hat{\delta}_j$ are obtained from the complete model M^{mdi} , see the Appendix A2 for estimation details. To construct the log-odds for the model containing only the macros (M^m), for example, we use (4.5) with $\hat{\gamma}_j$ and $\hat{\delta}_j$ set to zero. For the model with macros and common frailty (M^{md}), only $\hat{\delta}_j$ is set to zero.

The constructed log-odds can be substituted in (4.8) 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

$$R^2(\theta) = 1 - \frac{KL(\theta_{max}, \theta)}{KL(\theta_{max}, \theta_{na})}. \quad (4.9)$$

Note that (4.9) 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 (4.9) as the proportional reduction in variation due to the inclusion of latent factors, see Cameron and Windmeijer (1997). As mentioned earlier, for the standard linear regression model (4.9) 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.

4.3 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 4.2 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 4.1 for a listing of macroeconomic and financial data. This data enters the analysis in the form of annual growth rates, see Figure 4.3 for time series plots.

Figure 4.2: 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 US 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.

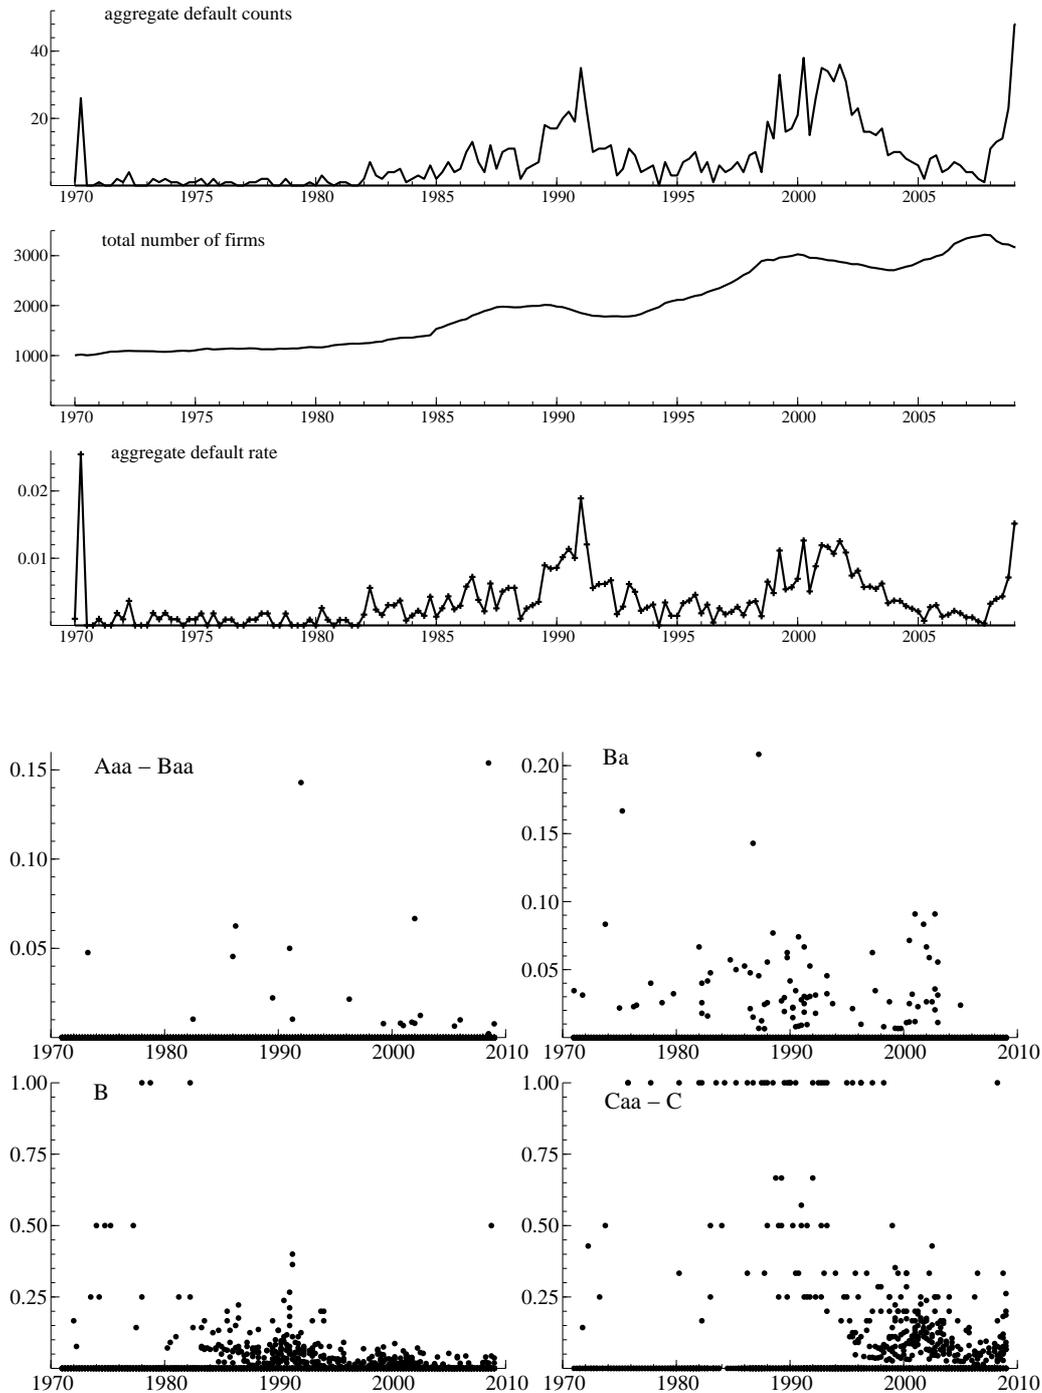


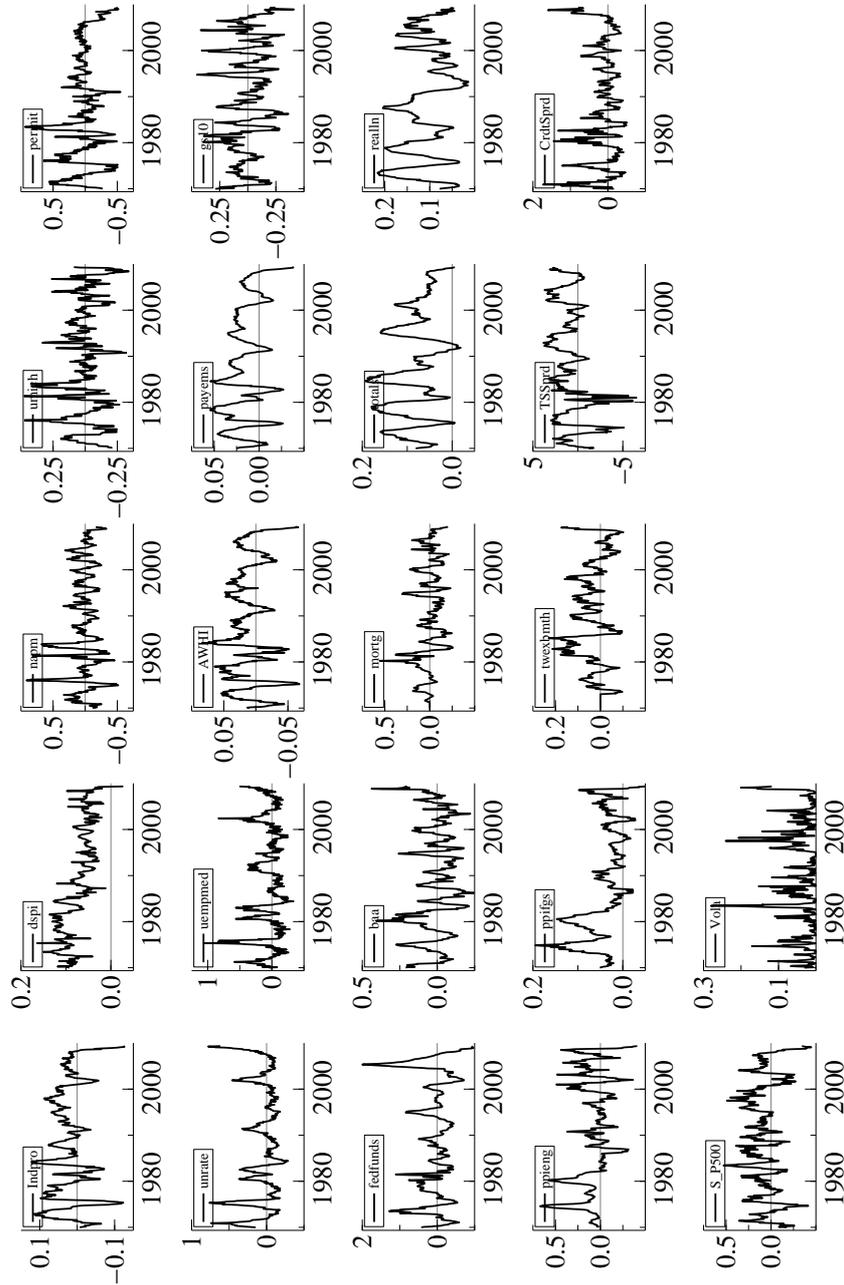
Table 4.1: Macroeconomic Time Series Data

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	5
	Disposable personal income	
	ISM Manufacturing index	
	Uni Michigan consumer sentiment	
	New housing permits	
(b) Labour market conditions	Civilian unemployment rate	4
	Median duration of unemployment	
	Average weekly hours index	
	Total non-farm payrolls	
(c) Monetary policy and financing conditions	Federal funds rate	6
	Moody's seasoned Baa corporate bond yield	
	Mortgage rates, 30 year	
	10 year treasury rate, constant maturity	
	Credit spread corporates over treasuries	
(d) Bank lending	Total Consumer Credit Outstanding	2
	Total Real Estate Loans, all banks	
(e) Cost of resources	PPI Fuels and related Energy	3
	PPI Finished Goods	
	Trade-weighted US dollar exchange rate	
(f) Stock market returns	S&P 500 yearly returns	2
	S&P 500 return volatility	

Figure 4.3: 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 4.1.



4.3.1 Major empirical results

Parameter estimates associated with the default counts are presented in Table 4.2. 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 $j = 1, \dots, J$ 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 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 4.2. 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 4.4 presents four estimated risk factors f_t^m as defined in (4.5) and (4.6). 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 4.1.

Figure 4.5 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 4.5 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 4.2.

Figure 4.6 presents smoothed estimates of the frailty and industry-specific factors. The

Table 4.2: 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 second column reports loading coefficients β_j on four common macro factors f_t^m . The third column reports the loading coefficients γ_j on the frailty factor f_t^d . The last column presents 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 4.4: Smoothed Macroeconomic Risk Factors

This figure presents four estimated risk factors f_t^m as defined in (4.5) and (4.6). We plot the estimated conditional mean of the factors, along with approximate standard error bands at a 95% confidence level.

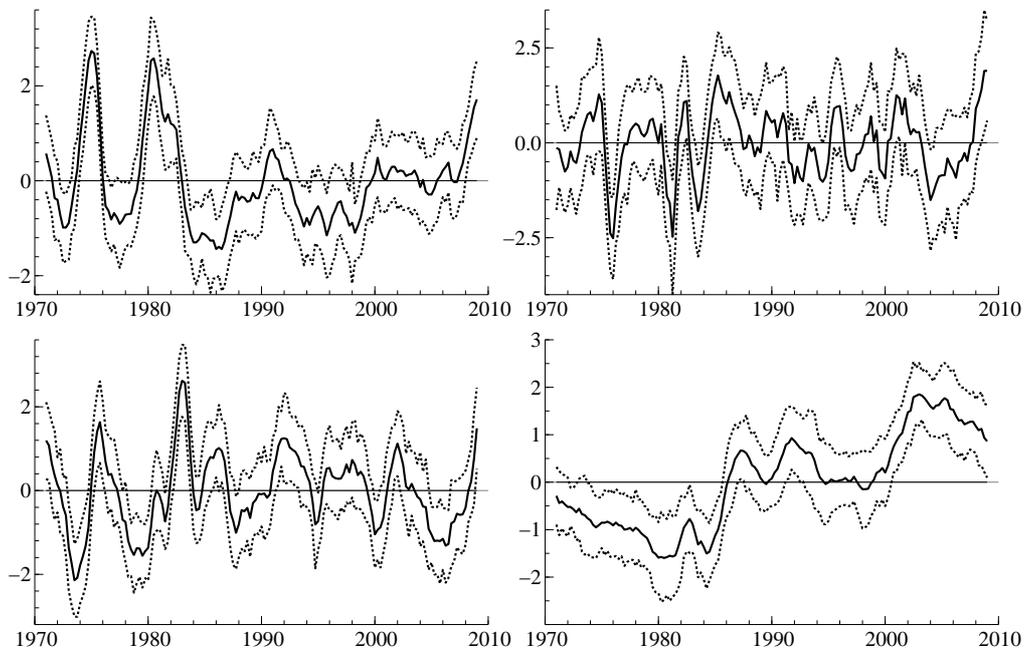


Figure 4.5: Shares of explained variation in macro and financial time series data

The figure indicates which share of variation in each time series listed in Table 4.1 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.

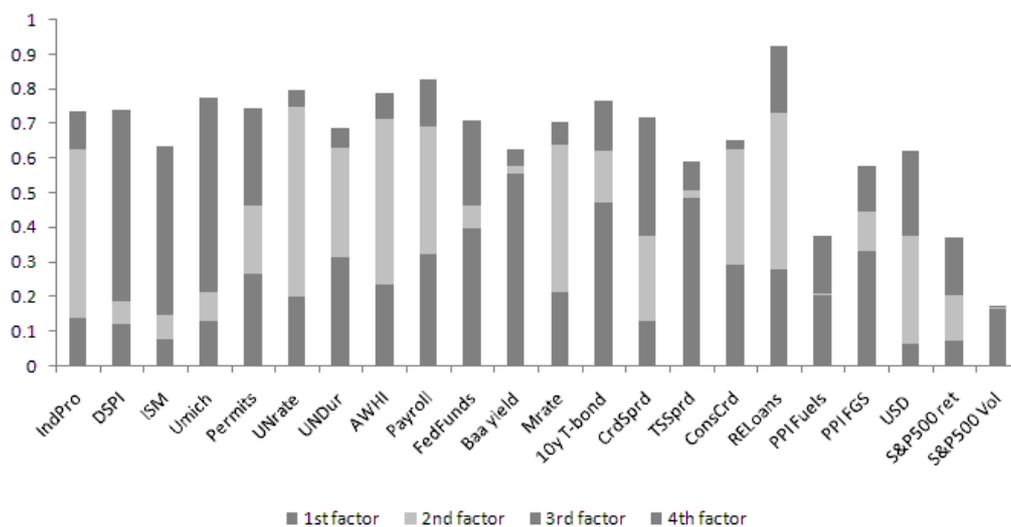
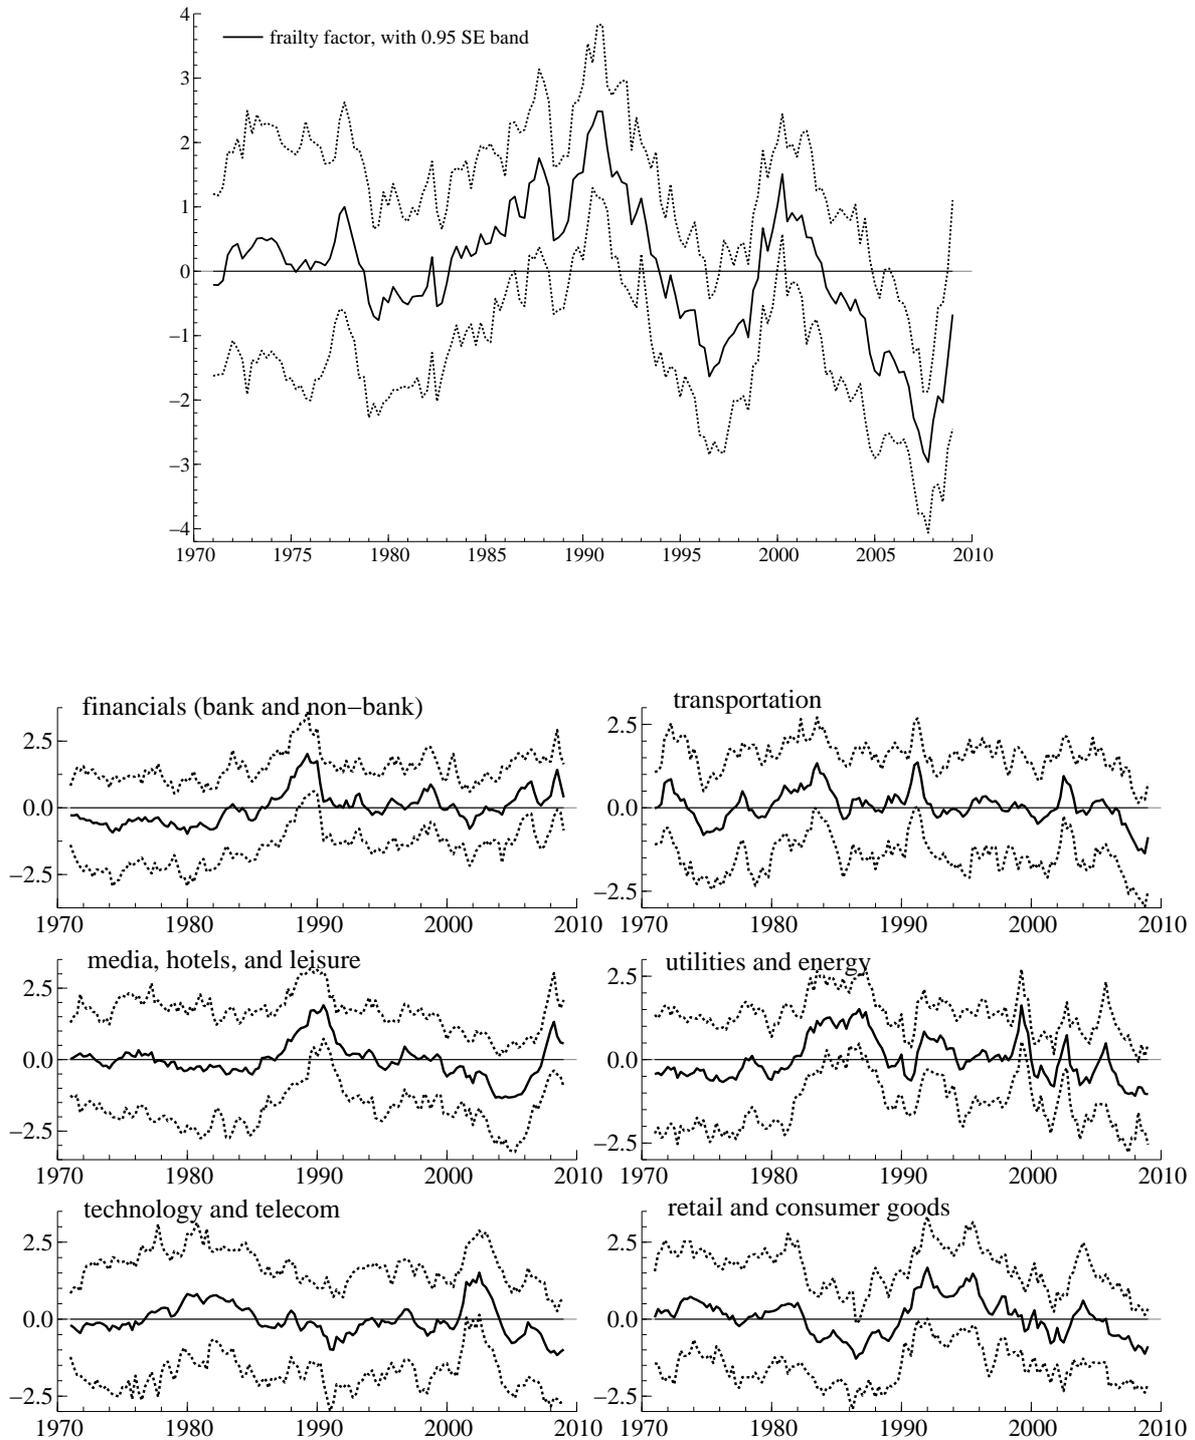


Figure 4.6: Smoothed Frailty Risk Factor and Industry-group dynamics

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.



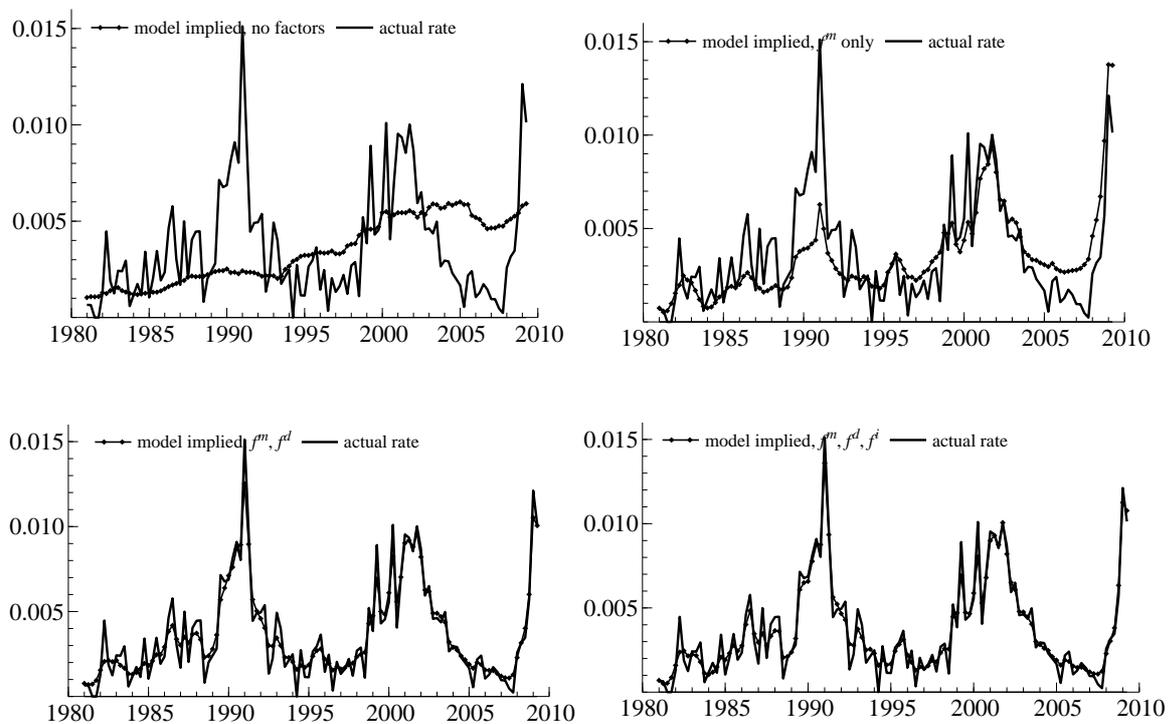
frailty factor is high *before* and *during* the recession years 1991 and 2001. As a result, the frailty factor implies additional default clustering in these times of stress. On the other hand, the large negative values before the 2007-2009 credit crisis imply defaults that are systematically ‘too low’ compared to what is implied by macroeconomic and financial data. 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 4.6 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 US 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 4.7 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

Figure 4.7: 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.



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 4.7 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.

4.3.2 Total default risk: a decomposition

We use the pseudo- R^2 measure as explained in Section 4.2.2 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 4.3 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 (26–53%). 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 4.3 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

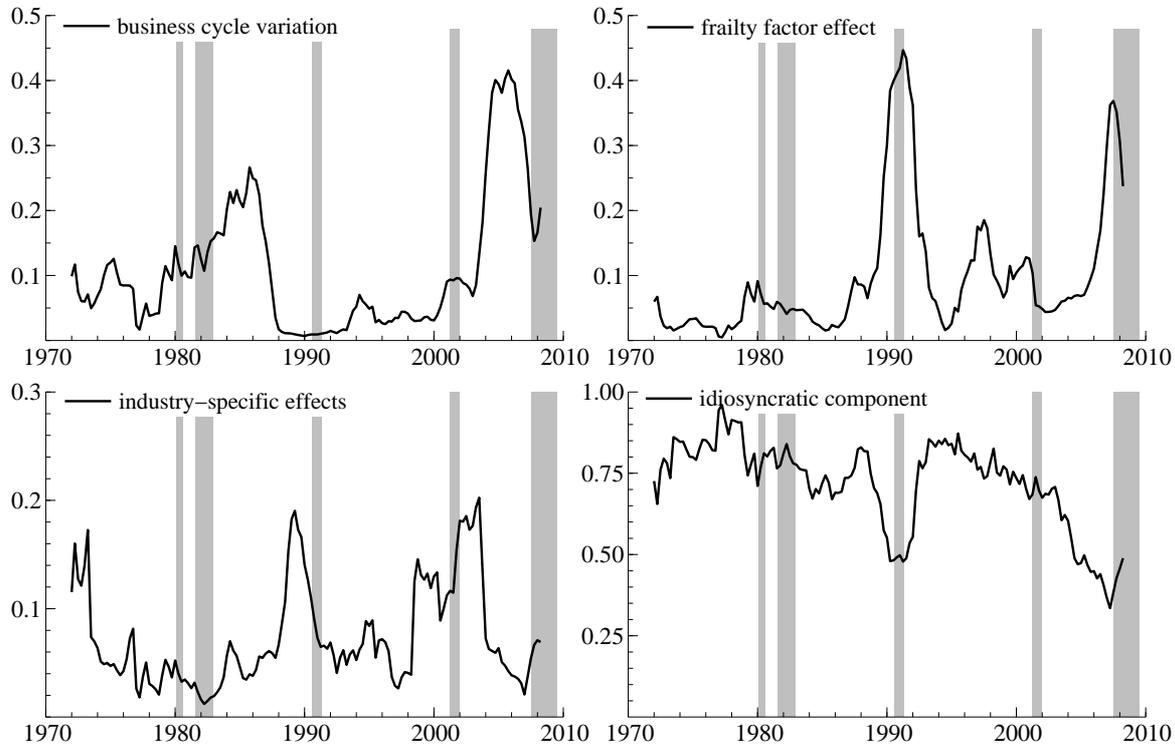
Table 4.3: A decomposition of total default risk

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 <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%
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 4.8: 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.



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 4.3 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 4.2.

Figure 4.8 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 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 4.8 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.

4.4 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 4.3.2 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.

4.4.1 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 4.3. 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 US 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 4.3 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 (4.2) directly to simulate future risk factor realizations. Finally, conditional on the risk factor path, the defaults can be simulated by combining (4.3) and (4.5).

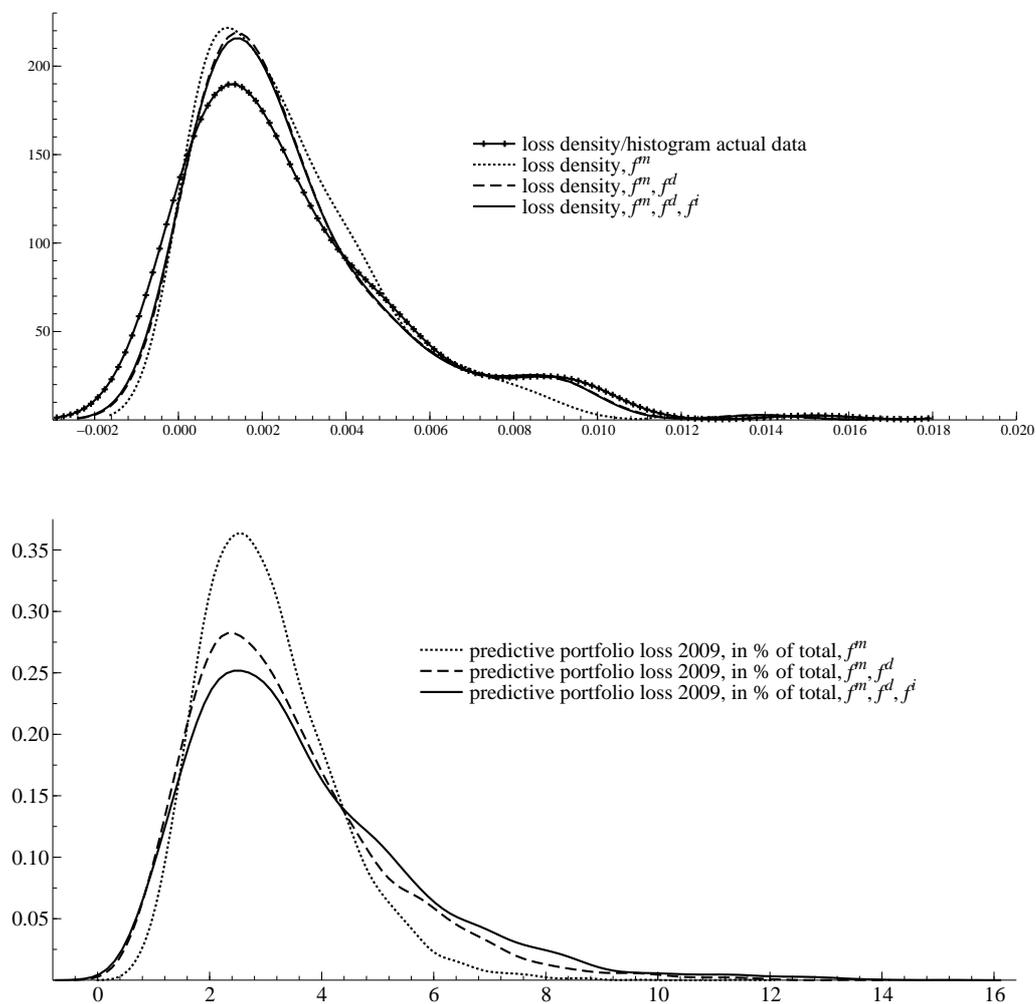
Our example portfolio is stylized in many regards. Nevertheless, it allows us to investigate the importance of macroeconomic, frailty, and industry-specific dynamics for the risk measurement of a diversified loan or bond portfolio.

The top panel in Figure 4.9 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 4.2. 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 4.9.

Figure 4.9: 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



By comparing the unconditional loss distributions in the top panel of Figure 4.9, 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 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.

4.4.2 Industry specific risk dynamics

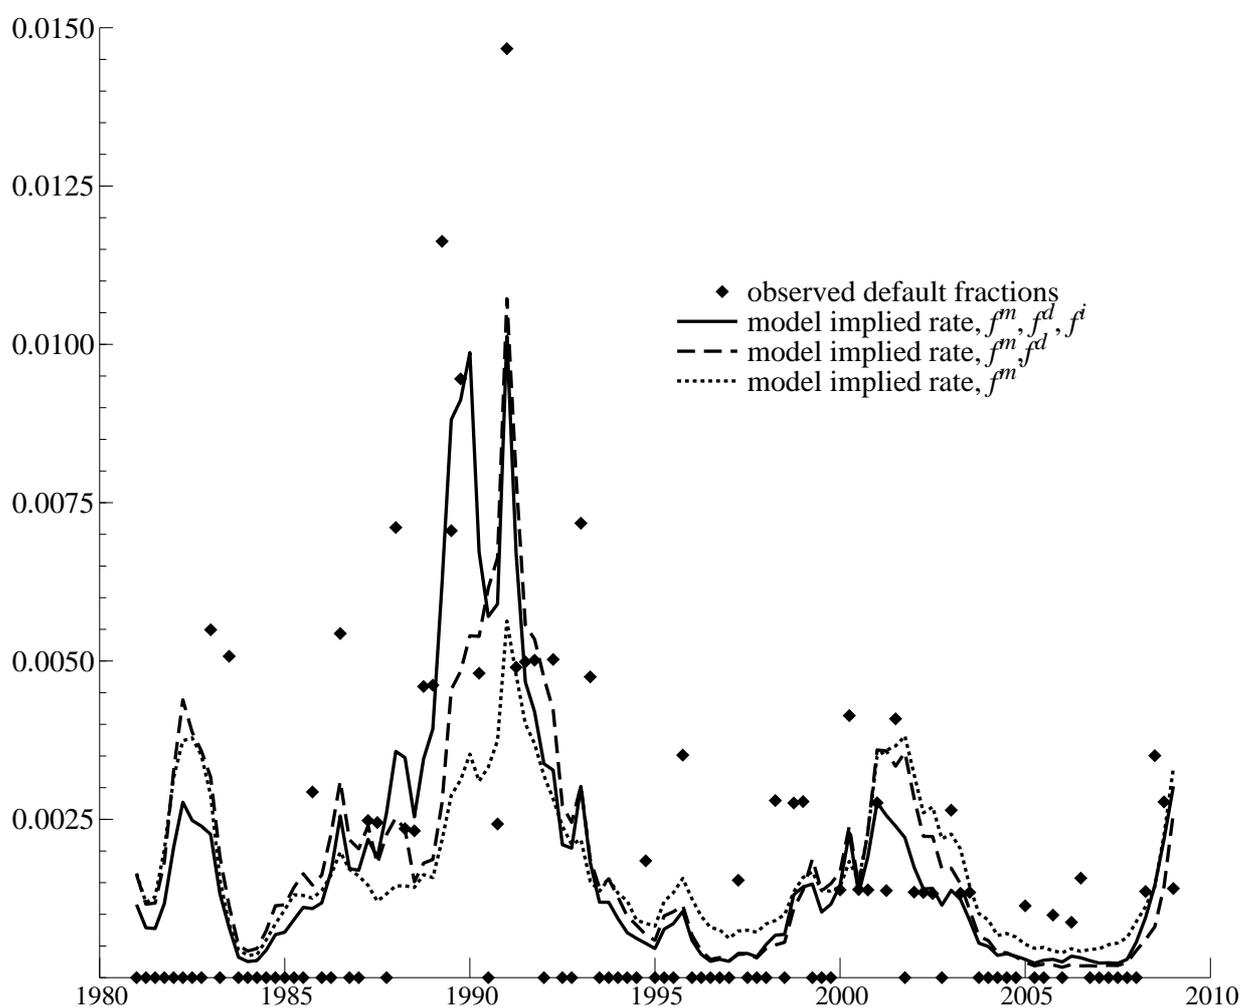
Section 4.3.2 shows that industry-specific variation accounts for about 17–37% of default rate volatility at the rating and industry level. Industry-specific factors capture the differential impact of each crisis on a given sector. For example, default stress for the banking industry has been high before 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 4.10 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

Figure 4.10: 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.



4.10. 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 4.6, captures the additional sector-specific stress during the banking crisis periods of 1986-1990 and 2007-09. It also 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 4.9 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.

4.5 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 the empirical study for US data, we found that approximately a third of default rate volatility at the industry and rating level is systematic. Approximately 30-60% of systematic risk variation is caused by macroeconomic and financial activity, while the remaining share is caused by frailty and industry factors, in roughly equal proportions. This suggests 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 about half of the systematic risk. This could be detrimental from a financial stability perspective.

We have given further empirical evidence that the composition of systematic risk varies over time. In particular, there has been 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 comovements between credits exposures increases through time, the fragility of the financial system may increase and prompt for an adequate (re)action.

Finally, the contributions of the frailty and industry components to systematic credit

risk volatility change over time. The effects turned out to be most important before and during high-default years in our sample. The direction of the effect, however, differs over time. Whereas frailty helps to capture rapid rises in defaults around the 1991 and 2001 recessions, the component is needed to explain the low default frequencies before the 2007-2008 financial crisis. During the more recent crisis, large changes in macro indicators suffice to pick up the rises in defaults over 2008.

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, unknown sources of default volatility such as frailty and contagion are needed to capture the patterns in the empirical data.

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

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

and we assume that y_{jt} given θ_{jt} is mutually and serially independent for $j = 1, \dots, J$ and $t = 1, \dots, T$. In other words, the dependence in the data set for our MiMe DFM is specified only via the signal θ_{jt} as given by (4.4) and (4.2). It follows that $p_j(y_{jt}|\theta_{jt}; \psi) = p_j(y_{jt}|\mathcal{F}_t; \psi)$ in (4.3).

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

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

The binomial density is obtained by having

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

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 (4.7) 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 $y = (y'_1, \dots, y'_T)'$ can be expressed by the joint density of y and

$f = (f'_1, \dots, f'_T)'$ where f is integrated out, that is

$$p(y; \psi) = \int p(y, f; \psi) df = \int p(y|f; \psi)p(f; \psi)df, \quad (\text{A.13})$$

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

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

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

$$p(y; \psi) = \int \frac{p(y, f; \psi)}{g(f|y; \psi)} g(f|y; \psi) df = g(y; \psi) \int \frac{p(y|f; \psi)}{g(y|f; \psi)} g(f|y; \psi) df. \quad (\text{A.14})$$

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

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

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

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

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}^2 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 (4.3), density $g(y|f; \psi)$ is based on the Gaussian density for $y_{jt} - \mu_{jt} - \theta_{jt} \sim N(0, \sigma_{jt}^2)$ (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

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

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

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

with \tilde{f}_t the t th 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 $g(f|y; \psi)$.

Chapter 5

A diagnostic framework for financial systemic risk assessment

5.1 Introduction

The debate on macro-prudential policies ignited by the recent financial crisis is currently under full swing. Various commentators agree that policy makers have heretofore overlooked important aspects of the macro-financial economic environment. For example, regulators have learned the hard way that asset and credit exposures that are safe on the micro level can have severe consequences if they are highly correlated (of high systematic risk) in the cross section. Such dependence undermines positive effects from diversification at the firm level, and may further lead to a ‘fallacy of composition’ at the systemic level, as described e.g. in Brunnermeier, Crocket, Goodhart, Persaud, and Shin (2009). Essentially, traditional risk-based capital regulation alone may underestimate systemic risk by neglecting the macro impact of banks reacting in unison to a shock.

Macro-prudential oversight seeks to focus on the financial system as a whole. It has recently been assigned to new regulatory councils, such as the European Systemic Risk Board (ESRB) in the European Union, and the Financial Services Oversight Council (FSOC) in the United States. Each board has a mandate to monitor, assess, and mitigate system-wide risk. There is widespread agreement that financial systemic risk is characterized by both cross-sectional and time-related dimensions, see e.g. ECB (2009). The cross sectional dimension concerns how risks are correlated across financial institutions at a given point in time (due to e.g. contagion across institutions, and prevailing default conditions at that time), while the time-series dimension concerns how systemic risk evolves over time (e.g. due to changes in the default cycle, changes in financial market conditions, and the gradual buildup of financial imbalances).

While there is broad consensus on the set of models, indicators, and analytical tools for e.g. macroeconomic and monetary policy analysis, such agreement is yet absent for macroprudential policy analysis. This paper seeks to help fill this important gap. In particular, we make three contributions to the literature on financial systemic risk assessment.

First, we propose a large-scale econometric framework for the measurement of global macroeconomic and default risk conditions. Essentially, the model is a high-dimensional diagnostic tool that allows us to track the evolution and composition of credit risk conditions around the globe. It is clear that what can't be measured can't be managed. As a result, a reliable and accurate diagnostic tool for systemic risk assessment is desperately needed. Among other issues when modeling credit market developments, commentators of the recent financial crisis stress the importance of an international perspective, see e.g. de Larosiere (2009). This requires looking beyond domestic developments to detect risks to financial stability. For the recent crisis, the saving behavior of Asian countries has been cited as a contributing factor to low interest rates and easy credit access in the U.S., see e.g. Brunnermeier (2009). Similarly, it is the exposure to developments in the U.S. housing market that triggered financial distress for European and Asian financial institutions. We therefore formulate and estimate a model for joint macroeconomic and credit risk conditions for firms in four broad geographical regions, i.e., (i) the U.S., (ii) current EU-27 countries, (iii) Asian countries, and (iv) all remaining countries. The large-dimensional model allows for the differential impact of world business cycle conditions on regional default rates, latent regional risk factors, as well as world-wide industry sector (contagion) dynamics.

The accurate measurement of point-in-time default conditions is complicated since not all processes that determine corporate distress are easily observed. Recent research indicates that observed macroeconomic and financial variables and firm-level information may not be sufficient to capture the large degree of default clustering present in observed corporate default data. In an important study for U.S. data, Das, Duffie, Kapadia, and Saita (2007) reject the joint hypothesis of well-specified default intensities in terms of (a) observed macroeconomic variables and firm-specific risk factors and (b) the conditional independence (doubly stochastic default times) assumption. In particular, there is substantial evidence for an additional dynamic unobserved 'frailty' risk factor, see McNeil and Wendin (2007), Azizpour and Giesecke (2008), Koopman, Lucas, and Monteiro (2008), Koopman and Lucas (2008), and Duffie, Eckner, Horel, and Saita (2009). 'Frailty' risk causes default dependence above and beyond what is implied by observed macroeconomic and financial data alone. For U.S. data, Koopman, Lucas, and Schwaab (2010) find that frailty effects account for 30–60% of the default rate volatility at the ratings and indus-

try level. The monitoring of unobserved frailty dynamics is thus of key importance in the measurement of credit risk conditions. Our model setup and suggested estimation methods accommodate such latent dynamic components.

As the paper's second main contribution, we present an indicator for latent common default stress for a given set of firms. When applied to financial firms, the indicator is a straightforward measure for financial systemic risk. The indicator is based on the above international factor model for disaggregated defaults. That means that it allows us to summarize and compare macroeconomic, frailty, and industry-specific (contagion) effects in defaults simultaneously. Our indicator complements alternative but very different measures of financial stability, as e.g. derived in Hartmann, Straetmans, and de Vries (2004), who rely on extreme value theory and market prices of financial assets, Segoviano and Goodhart (2009), who focus on a small set of systemically relevant financial firms and apply copula methods, and Adrian and Brunnermeier (2009), whose CoVar-measure is based on quantile regressions for sorted equity portfolios of financial firms.

Finally, in an empirical analysis of worldwide credit data for more than 12.000 firms in four broad economic regions, we provide evidence that the magnitude of 'frailty' effects at a given point in time can serve as a warning signal for macro-prudential policy makers. Frailty effects are highest when aggregate default conditions (the 'default cycle') diverge significantly from aggregate macroeconomic conditions (the 'business cycle'). Historically, frailty effects have been pronounced during bad times, such as the savings and loan crisis in the U.S. leading up to the 1991 recession, or exceptionally good times, such as the years 2005-07 leading up to the recent financial crisis. In the latter years, default conditions are much more benign than would be expected from observed macro and financial data, see also Koopman, Lucas, and Schwaab (2008). In either case, macro-prudential policy makers should be aware of a possible decoupling of systematic default risk conditions from macroeconomic and financial market conditions. However, we also demonstrate that such decoupling may not be easy to detect in real time. In this regard, the monitoring of both business and default cycle conditions is key.

The remainder of this paper is structured as follows. Section 5.2 presents the international modeling framework and introduces the systematic credit risk indicator (SRI) for common default stress. Section 5.3 presents the data for the empirical study. Estimates of international default conditions and corresponding risk factor estimates are presented in Section 5.4. Warning signals based on estimated frailty effects are discussed in Section 5.5. Section 5.6 concludes.

5.2 The econometric framework

We consider default count data y_t and macroeconomic covariates x_t given by

$$y_t = (y_{1,1t}, \dots, y_{1,Jt}, \dots, y_{R,1t}, \dots, y_{R,Jt})' \quad (5.1)$$

$$x_t = (x_{1t}, \dots, x_{Nt})' \quad (5.2)$$

for $t = 1, \dots, T$, where $y_{r,t} = (y_{r,1t}, \dots, y_{r,Jt})$ are default counts for a given region $r = 1, \dots, R$, and a certain combination of rating and industry group $j = 1, \dots, J$. The covariates in x_t summarize global macroeconomic and financial market conditions. We assume that both the default counts and the macro data are driven by dynamic factors.

The model combines normally and non-normally (discrete, count) distributed observations x_t and y_t , respectively. Conditional on latent factors f_t and observed covariates c_t , the measurements are independent. In our specific case, we assume that conditional on f_t and c_t , defaults $y_{r,jt}$ are binomially distributed with $k_{r,jt}$ trials and time-varying probability $\pi_{r,jt}$. Here, $k_{r,jt}$ denotes the number of firms in a specific ratings and industry bucket j in region r at time t , while $\pi_{r,jt}$ is the time-varying probability of default. Elements of x_t conditional on f_t follow a normal distribution with time-varying mean.

$$y_{r,jt} | f_t^m, f_t^d, f_t^i, c_t \sim \text{Binomial}(k_{r,jt}, \pi_{r,jt}) \quad (5.3)$$

$$x_{nt} | f_t^m \sim \text{Gaussian}(\mu_{nt}, \sigma_n^2), \quad (5.4)$$

where the factor structure distinguishes macro, frailty, industry-specific (contagion) effects, denoted f_t^m , f_t^d , f_t^i , respectively. The factors f_t^m capture shared business cycle dynamics in both macro data and default counts. Therefore, factors f_t^m are common to all data y_t and x_t . Frailty factors f_t^d are region-specific, i.e., they only load on defaults $y_{r,t} = (y_{r,1t}, \dots, y_{r,Jt})$ from a given region r . The frailty factors are independent of observed macroeconomic and financial data, and thus pick up any default-specific variation above and beyond that is implied by macro factors f_t^m . Latent factors f_t^i 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 capture the industry-specific propagation of macro and financial shocks. Finally, observed covariates c_t may include trailing default rates in the rest of the world in excess of respective domestic rates, and trailing industry-level default rates in excess of aggregate rates. As a result, default conditions ‘elsewhere’ can influence conditions ‘at home’ directly. Including other observed risk factors is also possible.

The point-in-time default probabilities $\pi_{r,jt}$ in (5.3) can also be referred to as (discrete time) hazard rates, or (discrete time) default intensities. They vary over time due to shared exposure to observed risk factors x_t as summarized by f_t^m , frailty effects f_t^d , latent industry specific effects f_t^i , as well as observed effects c_t at that time. We model $\pi_{r,jt}$ as the logistic transform of an index function $\theta_{r,jt}$,

$$\pi_{r,jt} = (1 + e^{-\theta_{r,jt}})^{-1}, \quad (5.5)$$

where $\theta_{r,jt}$ may be interpreted as the log-odds or logit transform of $\pi_{r,jt}$. This transform yields convenient expressions and ensures that conditional probabilities $\pi_{r,jt}$ are in the unit interval.

The panel data dynamics in observed data (5.1) and (5.2) are captured by time-varying parameters, or ‘signals’, given by

$$\theta_{r,jt} = \lambda_{r,j} + \beta'_{r,j} f_t^m + \gamma'_{r,j} f_t^d + \delta'_{r,j} f_t^i + \kappa'_{r,j} c_t, \quad (5.6)$$

$$\mu_{r,nt} = c_{r,n} + \beta'_{r,n} f_t^m, \quad (5.7)$$

where $\lambda_{r,j}$ is a fixed effect, and risk factor sensitivities $\beta_{r,j}$, $\gamma_{r,j}$, $\delta_{r,j}$, and $\kappa_{r,j}$ refer to macro factors, frailty factors, industry-specific factors, and observed covariates, respectively. Fixed effects and factor loadings may differ across firms and regions. Since the cross-section is high-dimensional, we follow Koopman and Lucas (2008) in reducing the number of parameters by imposing the following additive structure

$$\bar{\chi}_{r,j} = \chi_0 + \chi_{1,d_j} + \chi_{2,s_j} + \chi_{3,r_j}, \quad \bar{\chi} = \lambda, \beta, \gamma, \delta, \kappa, \quad (5.8)$$

where χ_0 represents the baseline effect, $\chi_{1,d}$ is the industry-specific deviation, $\chi_{2,s}$ is the deviation related to rating group, $\chi_{3,r}$ is the deviation related to economic region. Due to the common effect χ_0 , some specific coefficients need to be set to zero for model identification. This parameter specification combines model parsimony with a sufficiently large degree of flexibility to fit the data.

All latent factors are stacked into the vector $f_t = (f_t^m, f_t^d, f_t^i)'$, and follow simple autoregressive dynamics,

$$f_t = \Phi f_{t-1} + \eta_t, \quad \eta_t \sim \text{NID}(0, \Sigma_\eta) \quad (5.9)$$

where the coefficient matrix Φ is diagonal, and Σ_η is a covariance matrix of full rank.

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 business cycle dynamics in macro and default data. Similarly, the credit climate and industry default conditions are captured by persistent processes for f_t^d and f_t^i , such that they can capture the clustering of high-default years. The initial condition $f_1 \sim N(0, \Sigma_0)$ completes the specification of the factor process. The $m \times 1$ disturbance vectors η_t are serially uncorrelated. For identification, we require $\Sigma_\eta = I - \Phi\Phi'$. This implies $E[f_t] = 0$, $\text{Var}[f_t] = I$, and $\text{Cov}[f_t, f_{t-h}] = \Phi^h$ for $h = 1, 2, \dots$. It identifies loading coefficients $\beta_{r,j}$, $\gamma_{r,j}$, $\delta_{r,j}$, and $\kappa_{r,j}$ in (5.6) as risk factor volatilities (standard deviations) for firms in cross section (r, j) .

Equations (5.1) to (5.9) form a mixed measurement dynamic factor model (MM-DFM) as described in Koopman, Lucas, and Schwaab (2010) and Creal, Schwaab, Koopman, and Lucas (2010). Unfortunately, the model log-likelihood does not exist in closed form for this class of models, due to the combined presence of non-Gaussian data y_t and serially correlated unobserved factors f_t . We overcome this issue by casting the model (5.1) to (5.9) to state space form, and estimating parameters and latent factors using Monte Carlo maximum likelihood methods based on importance sampling. Missing values in both the default and macro data panel can be accommodated in a straightforward way. We refer to Creal et al. (2010) for estimation details, and to Durbin and Koopman (2001, part II) for a textbook treatment of non-Gaussian models in state space form.

5.2.1 Conditioning variables for capital requirements

There is mounting academic evidence that a macro-prudential policy maker can use variation in capital and liquidity requirements in the cross section and over time to mitigate systemic risk. In the *cross section*, macroprudential policies should differentiate between firms according to their systemic impact, see e.g. White (2004, 2006) for early proposals in that direction. Systemic stability is seen as a common good, and macroprudential policies seek to discourage strategies which cause a negative externality on the financial system. As a result, a large bank may be subject to more scrutiny and higher capital charges than, for example, a smaller insurance firm. Firm-specific measures of systemic risk, such as CoVar of Adrian and Brunnermeier (2009), seek to make this approach operational. There, institutional characteristics such as leverage, size, and maturity mismatch permit inference on systemic risk contribution. Perotti and Suarez (2009) argue for liquidity charges as a macroprudential tool in the cross section. Such liquidity charges are understood to be complementary to other reforms, such as counter-cyclical capital charges. In

our framework of Section 5.2, inference on systemic risk contribution could be read off the estimated risk factor sensitivities in (5.6), which are allowed to vary in the cross section according to (5.8). In practice, however, a large amount of pooling is required during estimation due to the sparse nature of observed default data.

In the *time dimension*, Goodhart and Persaud (2008) stress that bank capital requirements should be counter-cyclical with respect to the credit cycle. Times of optimism are the best times for financial institutions to make provisions for bad times. Incentives from current regulation are pro-cyclical, such that banks may respond to lower estimates of risk by going after the marginal customer in late booms. Goodhart and Persaud (2008) suggest that the rate of change in total bank lending and asset (equity) prices in relevant sectors could be used as conditioning variables for counter-cyclical capital charges. Angelini, Neri, and Panetta (2010) investigate changes in output, outstanding credit, and equity values as possible conditioning variables in a macro DSGE model. Their results indicate that time-varying capital requirements can cooperate with monetary policy to achieve both low inflation and financial stability. In a similar setting, Angeloni and Faia (2009) find an optimal policy combination that includes mildly counter-cyclical capital ratios and a monetary policy rule that reacts to inflation and asset prices. Covas and Fujita (2009) derive and calibrate a general equilibrium macro model with counter- and pro-cyclical capital requirements for banks. They find that output volatility decreases (by about 25%) and household welfare increases (by about 1.7%) if the countercyclical regime is adopted. Repullo and Suarez (2009) show that cyclical adjustments in the confidence level underlying Basel II regulation can reduce its procyclical effects on the supply of credit without compromising banks' long-run solvency targets. Interestingly, they condition counter-cyclical capital adjustments on estimates of the latent default cycle.

The case for differential capital (or liquidity) charges in the cross section or time dimension does not come without criticism. Changing capital requirements in response to conditioning variables is a costly and risky business. When designing capital requirements that address systemic concerns, regulators must weigh the costs such requirements impose on banks during good times against the benefit of having more capital in the financial system when a crisis strikes, see the report of the Squam Lake Working Group on Financial Regulation (2009). Second, long term effects of such regulatory change, such as e.g. on economic growth, are unknown. Other issues, such as 'rules vs. discretion' or Lucas critique arguments, are still open to discussion.

Despite the ongoing discussion about policy instruments, commentators broadly agree that what can't be measured can't be managed. As a result, diagnostic tools are needed that permit an assessment of current levels and the composition of shared risk. It is

against this background that we propose a new indicator for latent common default stress in the next section.

5.2.2 Systematic default risk indicator

This section presents an indicator for point-in-time default stress that is common to a given selection of firms. For example, we are interested in the measurement of latent default stress underlying U.S. and European financial firms. The indicator is based on the above reduced form modeling framework for international default conditions. As a result, it summarizes the information from observed risk factors, latent effects, and contagion dynamics at the industry and international level.

Default signals $\theta_{r,jt}$ in (5.6) capture the level and time variation in the log-odds of conditional default probabilities $\pi_{r,jt}$ over time for firms in region r and rating and industry combination j . Signals $\theta_{r,jt}$ consist of two terms,

$$\theta_{r,jt} = [\lambda_{r,j}] + [\beta'_{r,j}f_t^m + \gamma'_{r,j}f_t^d + \delta'_{r,j}f_t^i + \kappa'_{r,j}c_t],$$

where, fixed effects $\lambda_{r,j}$ pin down the through-the-cycle default rate, while systematic factors f_t^m , f_t^d , f_t^i , and c_t jointly determine *point-in-time* default conditions. Risk factors f_t^m and f_t^d load on all firms in a given region, and therefore do not average out in the cross section. Effects due to f_t^i and c_t average out only to some extent in a large loan portfolio. The aggregation of signals $\theta_{r,jt}$ over cross sections j can be achieved by pooling the factor loadings for these firms. Where this is not appropriate, one may use the fact that default intensities $\pi_{r,jt}$ are additive, see Lando (2003, Chapter 5). This allows to compute aggregate intensities, and the corresponding log-odds ratios.

Signals $\theta_{r,jt}$ are Gaussian since all risk factors in f_t are Gaussian. Standardized signals $z_{r,jt}$ are unconditionally standard normal, and obtained as

$$z_{r,jt} = (\theta_{r,jt} - \lambda_{r,j}) / \sqrt{\text{Var}(\theta_{r,jt})},$$

where $\text{Var}(\theta_{r,jt}) = \beta'_{r,j}\beta_{r,j} + \gamma'_{r,j}\gamma_{r,j} + \delta'_{r,j}\delta_{r,j} + \kappa'_{r,j}\kappa_{r,j} \geq 0$ is the unconditional variance of $\theta_{r,jt}$. Our systematic credit risk indicator (SRI) for firms of type j in region r at time t is given by

$$\text{SRI}_{r,jt} = 100\bar{\Phi}(z_{r,jt}), \quad (5.10)$$

where $\bar{\Phi}(z)$ is the standard normal cumulative distribution function. Values of $\text{SRI}_{r,jt}$ lie between 0 and 100 by construction with uniform (unconditional) probability. Values

below 50 indicate less-than-average common default stress, while values above 50 suggest above-average common stress. Values below 25, say, are exceptionally benign, and values above 75 are indicative of substantial systematic stress. A measure of financial systemic stress is obtained when (5.10) is based on model-implied hazards rates for financial firms in a given region. The indicator then summarizes common financial distress due to shared latent macro, frailty, and financial industry-specific effects.

5.2.3 Magnitude of frailty effects

We employ two complementary approaches to assess the magnitude of frailty effects over time. First, and most straightforwardly, we inspect the estimated frailty risk factors directly. All latent risk factors have zero mean and unit variance by construction. As a result, factor realizations larger than 1.64 in absolute value can be seen as ‘significant’, and values larger than two can be seen as ‘extreme’.

Second, we follow Koopman, Lucas, and Schwaab (2010) by relying on improvements in a (pseudo) R-squared measure for non-Gaussian data. This allows us to assess the relative magnitudes of macro, frailty, industry (contagion), and idiosyncratic risk in defaults. Such a decomposition can be based on a rolling window, which permits an assessment of such effects over time. Due to the non-linearity in (5.5) and the presence of multiple risk factors in (5.6), the same absolute value of frailty can have different effects on defaults depending on the value of the other factors at that time. Estimated reductions in a pseudo R-squared measure take this effect into account. On the other hand, defaults need to be sufficiently numerous to allow for a rolling window. For simplicity, we focus on improvements in the McFadden R-squared statistic

$$R_{\text{MF}}^2(\hat{\theta}) = 1 - l(\hat{\theta})/l(\theta_0), \quad (5.11)$$

where $l(\hat{\theta})$ is the log-likelihood of the fitted model, and $l(\theta_0)$ is the log-likelihood of a model that contains intercept terms only. The McFadden R-squared resembles the standard R-squared from linear regression. The two measures coincide when the log-likelihood of the intercept model is given by a total sum of squares, and the log-likelihood of the fitted model is a sum of squared errors. Other pseudo R-squared measures can also be considered.

5.3 Data

This section introduces the international data used in the empirical study below. We use data from two main sources. First, a panel of macroeconomic and financial time series data is taken from Datastream with the aim to capture international business cycle and financial market conditions. Macroeconomic data is obtained for $R = 4$ economic regions, i.e., (i) the U.S., (ii) E.U. countries, (iii) Asian countries, and (iv) all remaining countries. Figure 5.1 provides a listing of the macro data. The panel is unbalanced due to missing data in the beginning of the sample for some countries. The macro variables enter the analysis as annual growth rates from 1980Q1 to 2009Q4.

A second dataset is constructed from default data from Moody's. The database contains rating transition histories and default dates for all rated firms (worldwide) from 1980Q1 to 2009Q4. This data allows to determine quarterly values for $y_{r,jt}$ and $k_{r,jt}$ in (5.3). The database distinguishes 12 industry sectors which we pool into $D = 7$ industry groups: banks and financial non-banks (fin); transport and aviation (tra); hotels, leisure, and media (lei); utilities and energy (utl); industrials (ind); technology and telecom (tec); and consumer goods and retailing (ret). We consider four broad rating groups, investment grade $Aaa - Baa$, and three speculative grade groups Ba , B , and $Caa - C$. This yields a total of $J = 4 \times 7 \times 4 = 112$ different groups. The default panel is unbalanced, because not all combinations of region, industry sector, and rating group contain firms at risk (exposures) at all times.

Table 5.2 provides an overview of the international exposure and default count data. Corporate data is most numerous for the U.S., with E.U. countries second. Most firms are either from the industrial or financial sector, while firms from the the transportation or retailing sector are relatively less frequent. The bottom of Table 5.2 suggests that about 60% of all worldwide ratings are investment grade. European and Asian firms are more likely to be rated investment grade, with shares of 83% and 75%, respectively.

Table 5.3 lists the countries contained in each region, along with the respective number of defaults and firms. When counting exposures and defaults, a previous rating withdrawal is ignored if it is followed by a later default. If there are multiple defaults per firm, we consider only the first event. In addition, defaults that are due to a parent-subsidiary relationship are excluded. Such defaults typically share the same default date, resolution date, and legal bankruptcy date in the database. Inspection of the default history (text) and parent number confirms the exclusion of these cases. Figure 5.1 plots aggregate default counts, exposures, and observed fractions over time for each of the four regions.

The default data allows to construct observed contagion factors at the industry level,

Table 5.1: International macroeconomic time series data

We list the variables contained in the macroeconomic panel. The time series data enters the analysis as yearly (yoy) growth rates. The sample is from 1980Q1 to 2009Q4.

Region	Summary of time series in category	Total no
(i) United States	Real GDP Industrial Production Index Inflation (implicit GDP price deflator) Dow Jones Industrials Share Price Index Unemployment Rate, 16 years and older US Treasury Bond Yield, 20 years US T-Bill Yield, 3 months ISM Purchasing Managers Index	8
(ii) E.U. countries	Euro Area (EA16) Real GDP Euro Area (EA16) Industrial Production Index Euro Area (EA16) Inflation (Harmonized CPI) Euro Share Price Index, Datastram Euro Area (EA16) Unemployment Rate Euro Area (EA16) Gov't Bond Yield, 10 years Euro Interbank Offered Rate (Eruibor), 3 months Euro Area (EA16) Industrial Confidence Indicator	8
(iii) Asian countries	Japan: Real GDP Japan: Unemployment Rate Japan: Tokio Stock Exchange Index (Topix) China: Real GDP China: Shanghai Composite Share Price Index Russia: Real GDP Russia: Unemployment Rate India: Industrial Production	8
(iv) Other countries	Canada: Industrial Production Canada: Unemployment Rate (Metropolitan Areas) Brazil: Industrial Production Brazil: Unemployment Rate	4

Table 5.2: International default and exposure counts

The top panel presents default counts disaggregated across industry sectors and economic region. The middle panel presents the total number of firms counted from 1981Q1 to 2009Q4. The bottom panel presents the cross section of firms at risk ('exposures') at point-in-time 2008Q1 according to rating group and economic region.

Defaults	U.S.	Europe	Asia	Other	Sum
Bank	41	8	9	13	71
Financial	84	4	8	6	102
Transport	90	17	1	7	115
Media	127	2	0	2	131
Leisure	97	9	1	14	121
Utilities	24	2	0	5	31
Energy	79	0	1	6	86
Industrial	435	16	16	37	504
Technology	177	38	3	21	239
Retail	94	1	2	2	99
Cons Goods	120	8	3	14	145
Misc	31	0	4	12	47
Sum	1399	105	48	139	1691

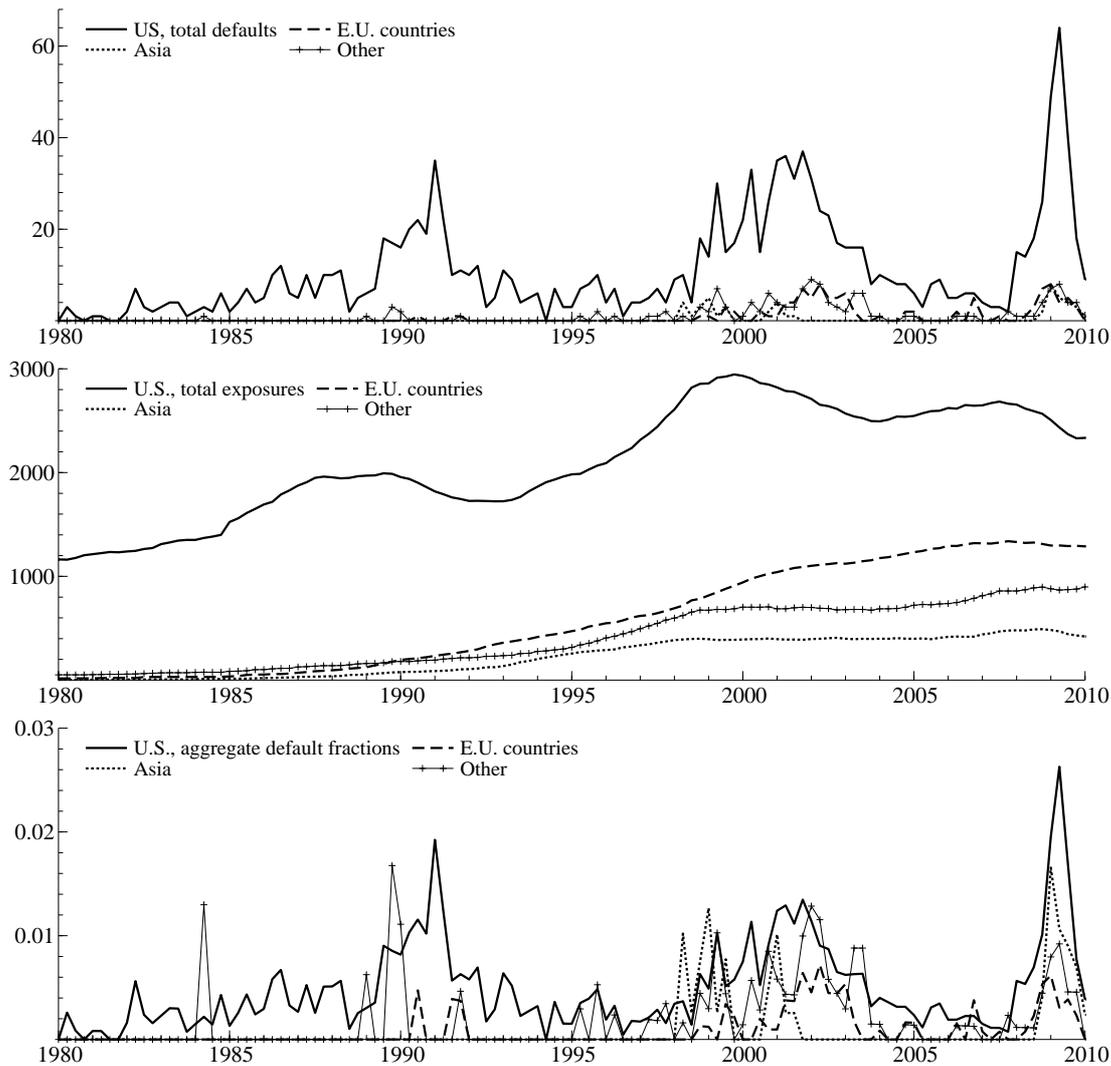
Firms	U.S.	Europe	Asia	Other	Sum
Bank	478	603	238	353	1672
Financial	966	371	130	370	1837
Transport	336	70	29	43	478
Media	460	33	5	29	527
Leisure	434	73	5	54	566
Utilities	597	149	41	97	884
Energy	512	84	31	121	748
Industrial	1920	419	180	317	2836
Technology	941	204	86	134	1365
Retail	311	32	21	25	389
Cons Goods	591	110	34	78	813
Misc	250	151	66	192	659
Sum	7796	2299	866	1813	12774

Firms, 2008Q1	U.S.	Europe	Asia	Other	Sum
Aaa	50	84	26	43	203
Aa	141	355	85	165	746
A	415	403	161	176	1155
Baa	575	229	91	200	1095
Ba	278	72	51	126	527
B	673	96	62	121	952
Caa-C	379	58	7	49	493
Sum	2511	1297	483	880	5171

in addition to latent industry effects. Default contagion is a possible alternative source of default clustering in observed data, see Jorion and Zhang (2007a), Lando and Nielsen (2008), and Boissay and Gropp (2010). We assume that default contagion due to business links is most important at the industry level. For example, a defaulting manufacturing firm may weaken other up- or downstream manufacturing firms. To capture industry-level (contagion) dynamics, we regress trailing one year default rates at the industry-level on a constant and the trailing one year aggregate default rate. Observed contagion factors are then obtained as the resulting standardized residuals. In this way, we eliminate the effect of the common macro and frailty factors and we retain industry-specific default variation. Additional latent industry-specific factors may capture remaining industry-level risk.

Figure 5.1: World default rates

We present time series plots of (a) the total default counts $\sum_j y_{r,jt}$ aggregated to a univariate series, (b) the total number of firms $\sum_j k_{r,jt}$ in the database, and (c) aggregate default fractions $\sum_j y_{r,jt} / \sum_j k_{r,jt}$ over time. We distinguish four different economic regions: the U.S., current E.U. countries, Asian countries, and remaining countries as listed in Table 5.3. The sample is from 1980Q1 to 2009Q4.



5.4 Major empirical results

This section presents the main empirical findings. Section 5.4.1 comments on parameter and risk factor estimates. Estimates of the latent default cycle and financial systemic distress are presented in Section 5.4.2.

5.4.1 Macro, frailty, and contagion effects

Table 5.4 presents the parameter estimates pertaining to model specification (5.1) to (5.9). The fixed effects and factor loadings in the signal equation (5.6) adhere to the additive structure (5.8). Our discussion focuses on the international aspects, since U.S. default data is already discussed in detail elsewhere, see e.g. Duffie et al. (2007), Duffie et al. (2009), and Koopman et al. (2010). Parameter estimates are presented in Table 5.4. Coefficients λ in the left column combine to baseline hazard rates. The middle and right columns present estimates for loadings β , γ , δ and κ pertaining to macro, frailty, industry (contagion), and observed trailing rates, respectively.

We find that macro, frailty, and industry-specific effects are all important for defaults. Defaults from all regions load significantly on common factors from macro and financial data f_t^m . Shared exposure to common macro factors already implies a considerable degree of default clustering. In general, however, common variation with macro data is not sufficient. Frailty effects are particularly pronounced for U.S. and Asian firms. The frailty factor loading for European firms is not significant. Latent industry-specific factors, however, load on all default data. As a result, latent dynamics are important for firms in each region.

Figure 5.2 presents the macro factor estimates. The factors are ordered from top left to bottom right according to their share of explained variation in macro data as listed in Table 5.1. They account for 19%, 18%, 17%, and 10% (64% in total) of international macro data variation. Factors 1 and 2 are most important for U.S. macro data, with a respective 27% and 19% share. Factors 2 and 3 load mostly on European data, with a respective 28% and 27% share. Factors 3 and 4 load heavily on Asian and the remaining countries, respectively. Contrasting this information with the middle column of Table 5.4 suggests that that macroeconomic developments elsewhere matter for domestic defaults. While the first macro factor loads mostly on U.S. macro data, it also loads heavily on European and Asian defaults, and more so than on U.S. defaults. This suggests that U.S. business cycle dynamics may matter *more* to European and Asian rated firms than to U.S. firms. The reverse is not the case. While factors 2 and 3 load heavily on European and Asian macro data, they do not load significantly on U.S. defaults. As a result, European and Asian macroeconomic developments matters mostly to rated firms that are also located in these regions. Ait-Sahalia, Cacho-Diaz, and Laeven (2010) find a similar asymmetry for international equity data.

Industry-specific dynamics are important for international defaults even after conditioning on common macro and frailty factors. Figure 5.3 presents estimates of scaled

Table 5.4: Parameter estimates

We report the maximum likelihood estimates of selected coefficients in the specification of the log-odds ratio (5.6) with parameterization (5.8) for λ and β . Coefficients λ combine to fixed effects, or baseline hazard rates. Factor loadings β , γ , δ , and κ refer to macro, frailty, industry (contagion), and observed risk factors, respectively. Monte Carlo log-likelihood evaluation is based on $M = 5000$ importance samples. The estimation sample is from 1981Q1 to 2009Q4.

Intercepts λ_j			Loadings f_t^m			Loadings f_t^d, f_t^i, c_t		
par	val	t-val	par	val	t-val	par	val	t-val
λ_0	-2.67	11.56	$\beta_{1,0}$	0.34	3.60	γ_{US}	0.52	6.71
			$\beta_{1,1,fin}$	0.02	0.11	γ_{EU}	-0.02	0.07
$\lambda_{1,fin}$	0.14	0.54	$\beta_{1,2,IG}$	0.78	2.18	γ_{AS}	1.52	1.91
$\lambda_{1,tra}$	-0.29	0.73	$\beta_{1,3,EU}$	0.26	0.92	γ_{SA}	0.49	2.03
$\lambda_{1,lei}$	-0.04	0.24	$\beta_{1,3,AS}$	1.02	1.38			
$\lambda_{1,utl}$	-0.27	0.84	$\beta_{1,3,Ot}$	0.38	1.45	δ_{fin}	0.45	2.35
$\lambda_{1,tec}$	-0.04	0.23				δ_{tra}	1.00	2.84
$\lambda_{1,ret}$	-0.40	2.49	$\beta_{2,0}$	0.21	1.75	δ_{lei}	0.34	1.54
			$\beta_{2,1,fin}$	0.08	0.53	δ_{utl}	0.99	2.34
$\lambda_{2,IG}$	-7.67	12.24	$\beta_{2,2,IG}$	-0.11	0.39	δ_{ind}	-	-
$\lambda_{2,BB}$	-3.78	10.03	$\beta_{2,3,EU}$	0.10	0.37	δ_{tec}	0.23	1.77
$\lambda_{2,B}$	-2.27	14.19	$\beta_{2,3,AS}$	0.01	0.01	δ_{ret}	0.46	2.28
			$\beta_{2,3,Ot}$	-0.08	0.30			
$\lambda_{3,EU}$	-1.23	5.45				δ_{+EU}	1.07	3.59
$\lambda_{3,AS}$	-1.80	3.00	$\beta_{3,0}$	-0.05	0.44	δ_{+AS}	-0.52	1.30
$\lambda_{3,Ot}$	-1.51	1.55	$\beta_{3,1,fin}$	-0.09	0.62	δ_{+Ot}	0.22	1.32
			$\beta_{3,2,IG}$	0.33	1.24			
			$\beta_{3,3,EU}$	-0.34	1.04	κ_{fin}	-0.05	0.37
			$\beta_{3,3,AS}$	0.13	0.18	κ_{tra}	-0.45	3.60
			$\beta_{3,3,Ot}$	-0.51	2.16	κ_{lei}	-0.10	1.27
						κ_{utl}	-0.20	1.54
			$\beta_{4,0}$	-0.11	1.44	κ_{ind}	-0.03	0.60
			$\beta_{4,1,fin}$	-0.05	0.36	κ_{tec}	0.25	3.37
			$\beta_{4,2,IG}$	0.92	3.65	κ_{ret}	0.13	1.44
			$\beta_{4,3,EU}$	-0.01	0.62			
			$\beta_{4,3,AS}$	-0.57	1.12	κ_{-US}	-0.00	0.07
			$\beta_{4,3,Ot}$	-0.24	1.44	κ_{-EU}	0.14	1.72
						κ_{-AS}	-0.14	0.49
						κ_{-OT}	0.04	0.32

Figure 5.2: Macro factors and explained variation

The top graph presents four common factors f_t^m from international macro and default data. The conditional mean estimates are plotted with corresponding 95% standard error bands. The bottom chart indicates which share of variation in each time series of Table 5.1 can be attributed to each macro factor in the top graph.

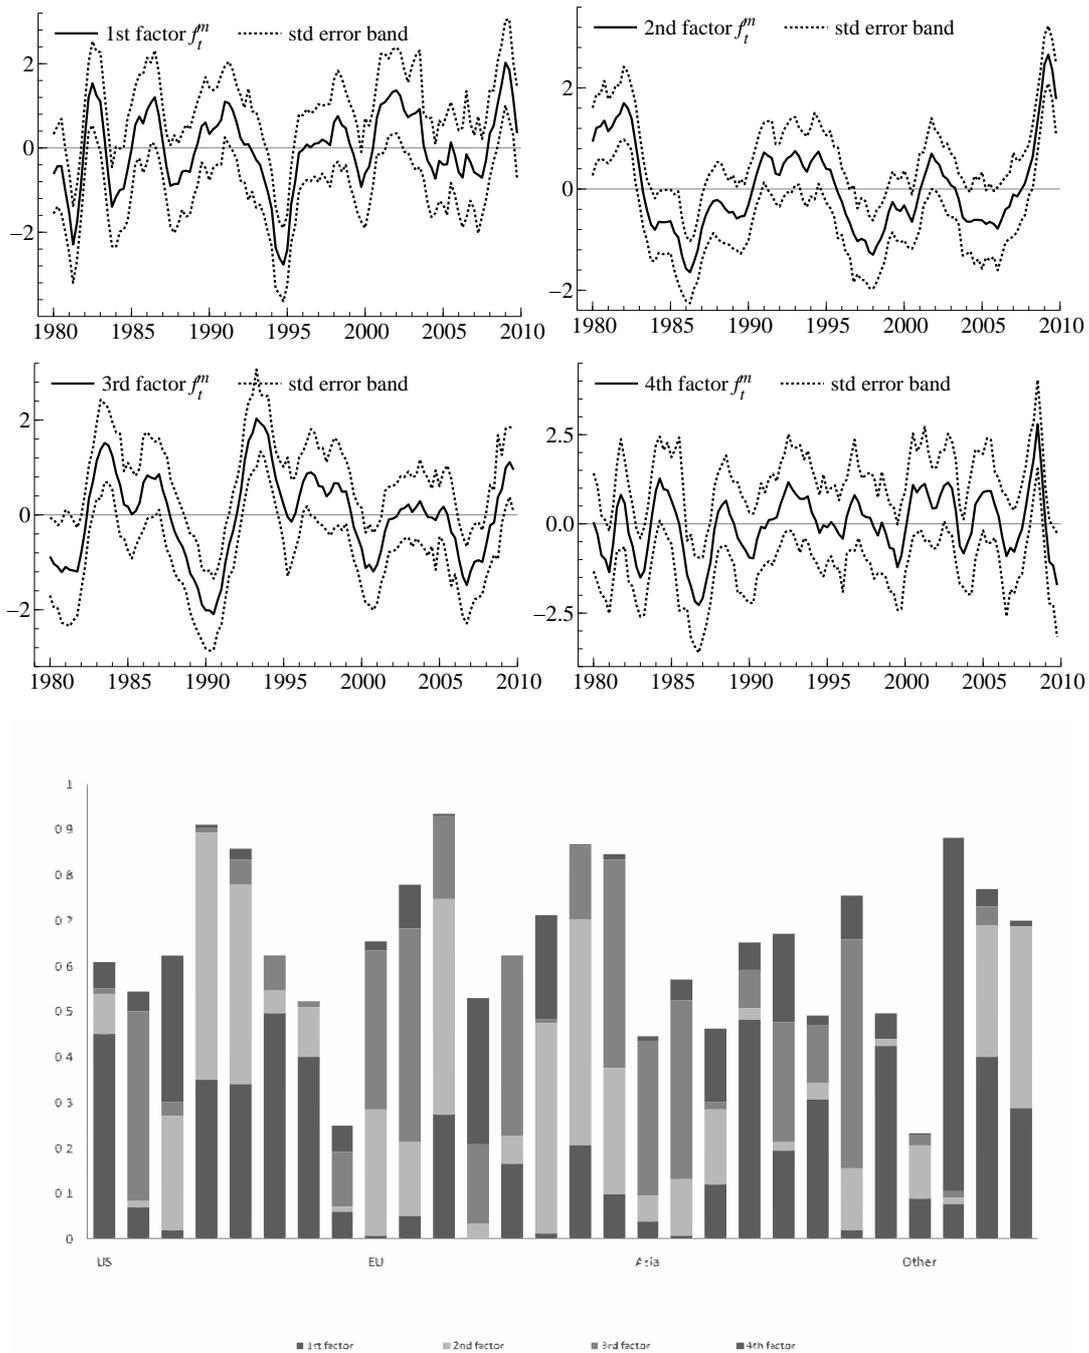


Figure 5.3: Industry-specific variation

The graph presents estimated industry-sector specific risk factors f_t^i . Factors are scaled by their volatility (loading) coefficient estimate for the U.S. We also indicate the share of total industry-specific risk that can be attributed to observed trailing industry-level default rates.

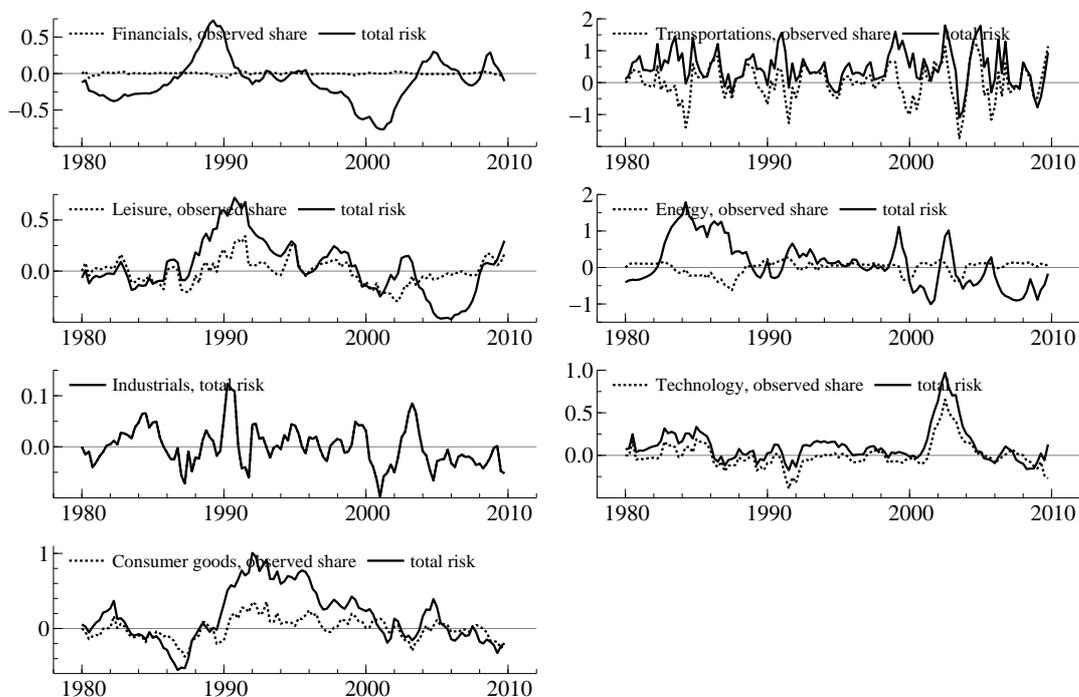
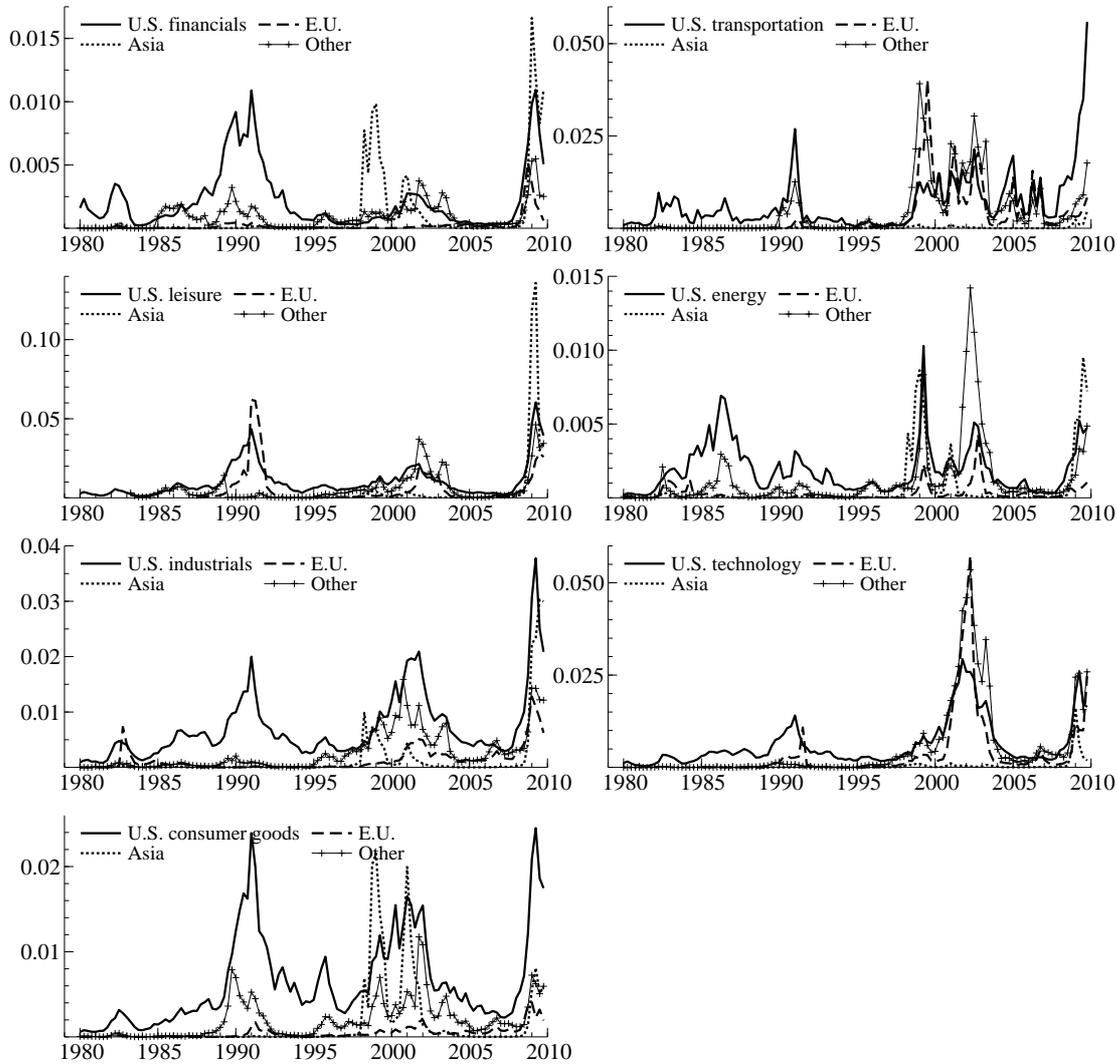


Figure 5.4: Default intensity estimates

We plot model-implied discrete time default intensities/hazard rates for seven broad industry groups and four economic regions. Estimation sample is 1981Q1 to 2009Q4.



industry-specific risk. Observed trailing industry-level default rates can approximate industry default dynamics only to some extent. Latent industry-specific factors still play an important role, in particular for European firms. The parameter estimates for δ in Table 5.4 suggest that European defaults load more heavily on industry-specific factors than U.S. and Asian firms.

Macro, frailty, and industry-specific factors combine to imply default rates at the regional and industry level. Figure 5.4 plots estimated default hazard rates for different regions and seven broad industry groups. Hazard rates tend to be lower for European

and Asian firms, reflecting higher credit ratings on average, see the bottom panel of Table 5.3. The recent financial crisis of 2008-09 is clearly visible in the financial, transportation, leisure, industrials, and consumer goods segments. Similarly, the default stress following the dot-com asset price bubble burst in 2000 can be seen in the technology sector from 2001-02, in each region. The model-implied default rate volatility at the industry level is striking. Point-in-time default hazard rates are five times (transportation, industrials) and up to ten times (financials) higher in bad times than in good times, in all regions.

5.4.2 Common default stress

The systematic risk indicator (5.10) allows to summarize estimated macro, frailty, and industry-specific effects when assessing common default stress for a given portfolio of firms. Figure 5.5 plots estimates of the default cycle for the U.S. and E.U. based on non-financial firms in the respective region. Shaded areas represent U.S. recession periods according to the NBER.

Not surprisingly, Figure 5.5 indicates that each U.S. recession period coincides with high levels of systematic default risk. In addition, the years around the 1991 and 2001 recession years suggest that a recession is not necessary for elevated default stress. Default conditions are particularly benign during the mid-1990s and mid-2000s, with values about 0.25 on a scale from zero to one. The mid-1990s are associated with the Clinton-Greenspan policy mix of low interest rates and low budget deficits, and corresponding favorable macroeconomic conditions. The mid-2000s are characterized by exceptionally low interest rates and easy credit access for U.S. firms. A comparison of U.S. and E.U. default conditions suggests that U.S. conditions are much more benign during the years 2005-07 leading up to the recent crisis. This may reflect direct and indirect effects from aggressive interest rate cuts by the Federal Reserve in these years, see e.g. Taylor (2008).

Figure 5.6 plots the model-implied world credit cycle using (5.10) and estimated default hazard rates across firms from all regions. Regional default cycle are plotted for comparison. The regional cycles are highly correlated, due to shared exposure to macroeconomic and industry-specific factors. The main peaks and troughs coincide. However, we (again) observe a pronounced dispersion of default conditions in the years 2005-07 leading up to the recent financial crisis. U.S. conditions are most benign, followed by Asian and 'other' countries, while European firms experience average to slightly above-average default stress. Whether this pronounced dispersion in default conditions can be related to monetary policy decisions during that time is an interesting question for future research.

Figure 5.5: Latent default cycles

The figure presents indicator values (5.10) based on estimated risk factors and parameter values for each region. The indicator captures default stress common to rated (non-financial) firms.

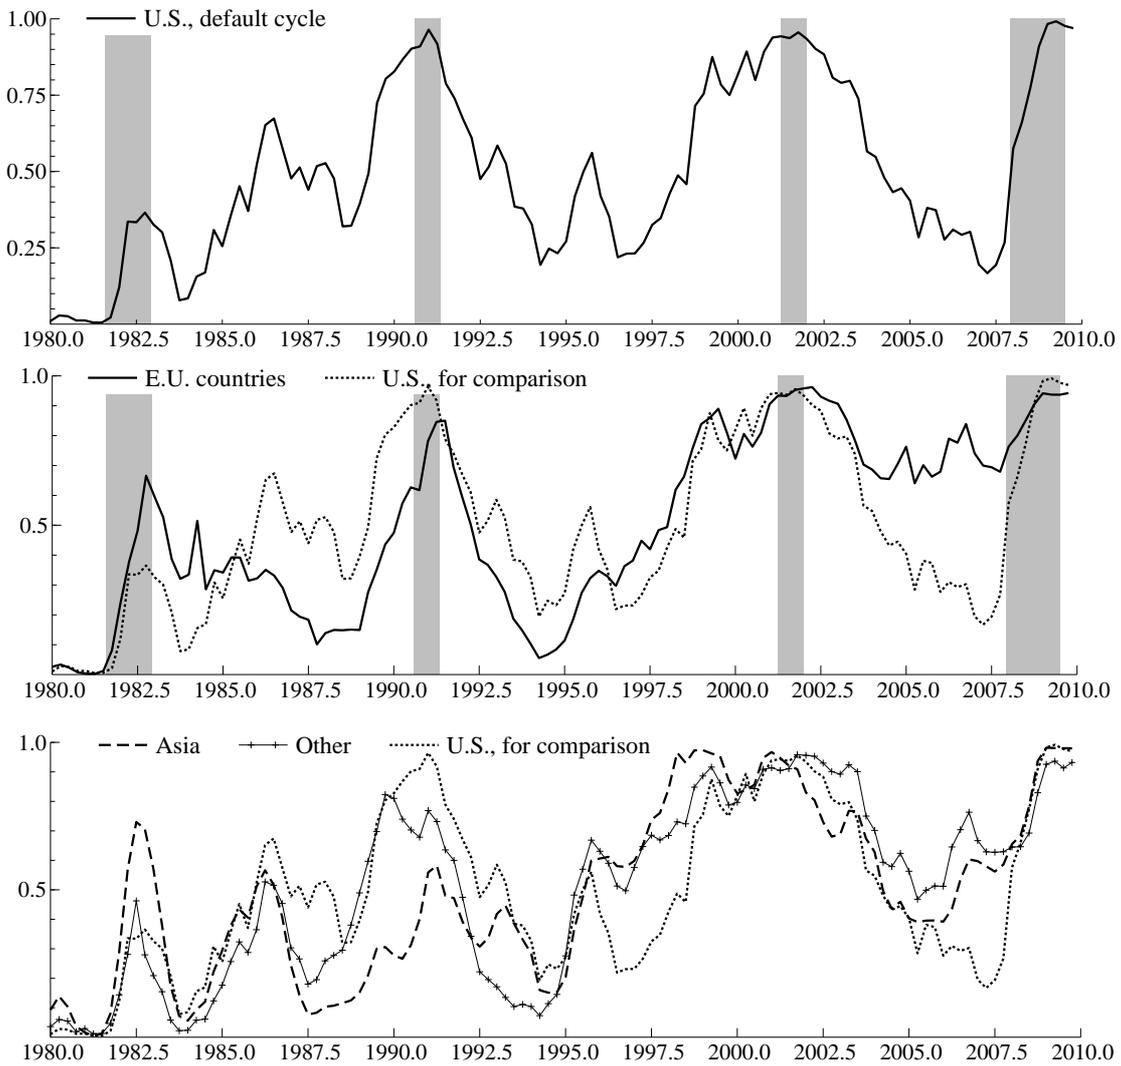


Figure 5.6: World credit cycle and regional deviations

We plot a world credit cycle estimate based on the indicator (5.10) and point-in-time default hazard rates for all (financial and non-financial) firms in each region. Regional credit cycle conditions are plotted for comparison.

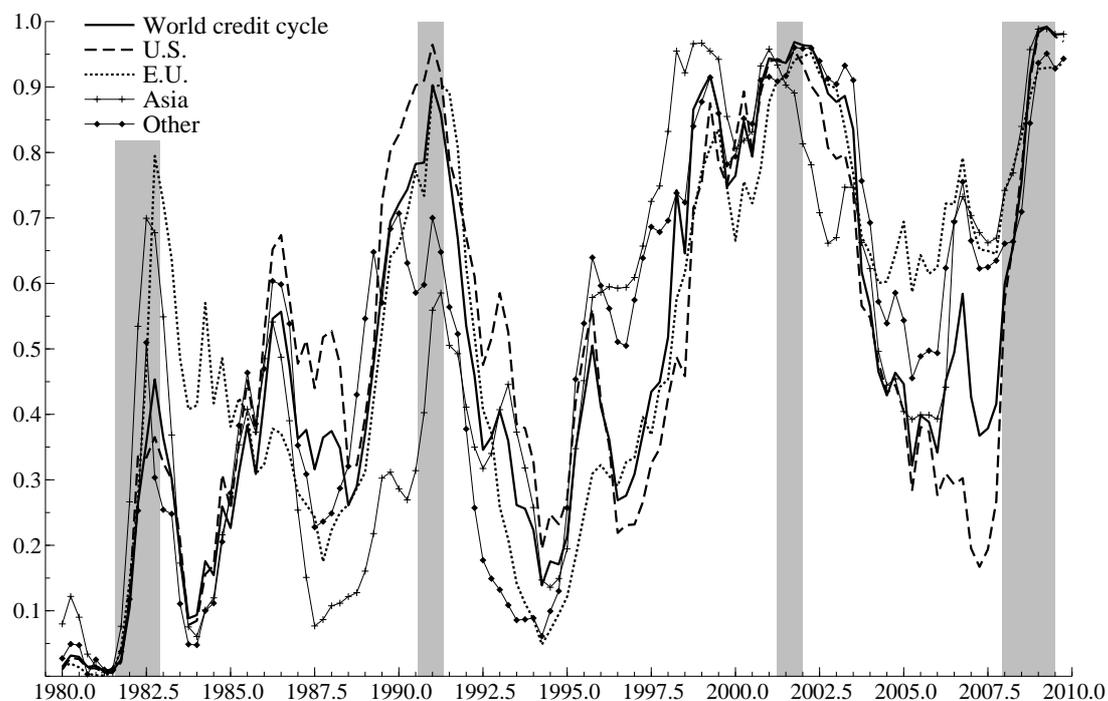


Figure 5.7: Common default stress for financials

We plot estimates of common default stress underlying financial firms (banks and financial non-banks) over time in different regions. Estimation sample is 1981Q1 to 2009Q4.

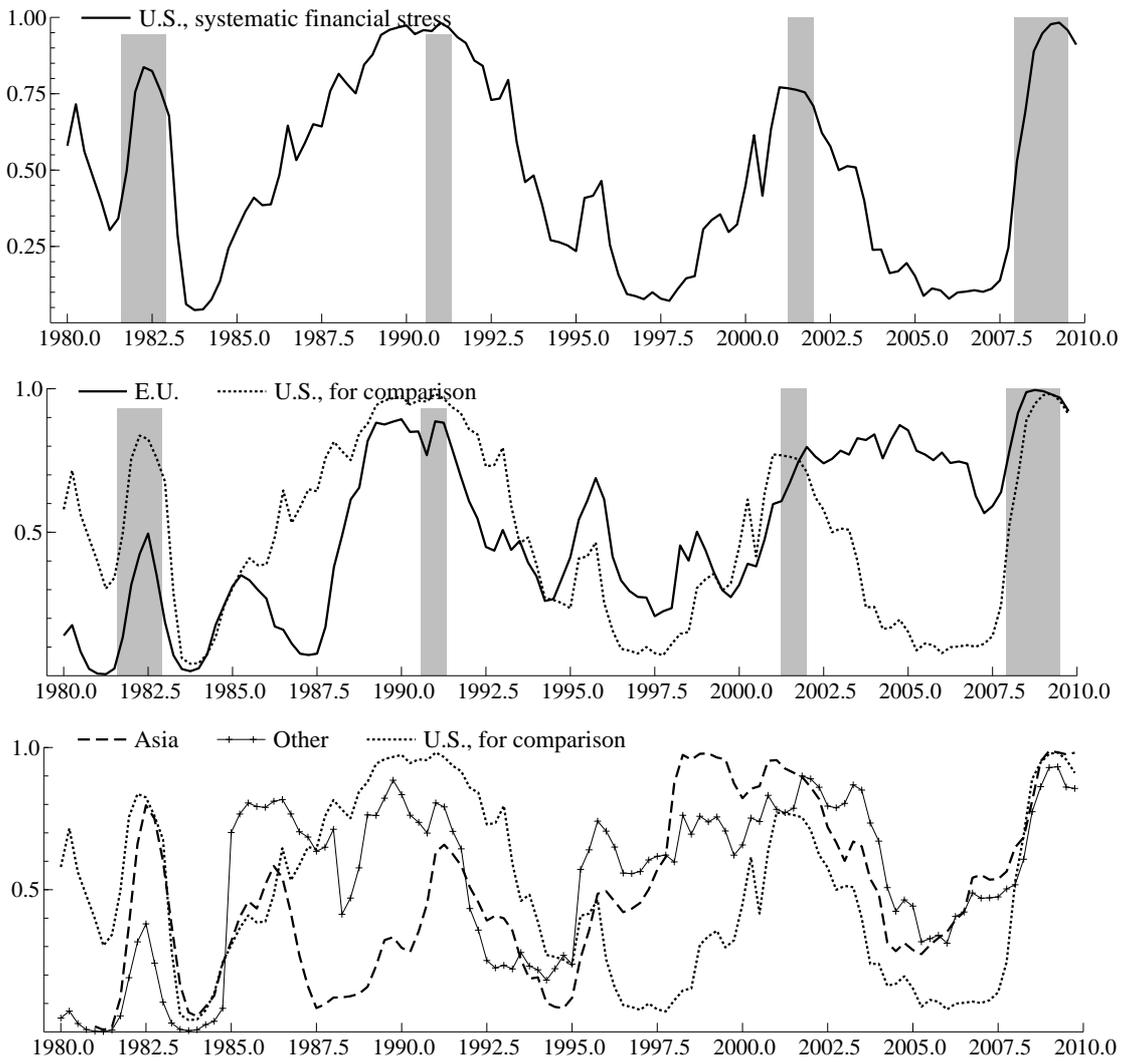


Figure 5.7 presents estimates of common distress for financial institutions based on the indicator (5.10). Shaded areas represent U.S. recession periods according to the NBER. Each recession implies common (and therefore systemic) stress on financials. The 1991 and 2008-09 recessions have been harder on U.S. financials than the relatively more benign 2001 recession. Furthermore, a recession is not necessary to have systemic stress on financial firms. An example for the latter are the late 1980s in the U.S., when common stress is pronounced while the economy is not in recession. European and Asian financial distress is correlated with, but different from, U.S. financial distress. For example, while default stress is virtually absent for U.S. financials during the mid-2000s, this is not the case for model-implied European conditions. Also, the mid-2000s do not appear to be particularly benign for European financials.

5.5 Early warning signals

Past experiences of financial fragility, financial booms and financial crisis, suggests that problems rarely appear at the same place in the financial system twice in a row. Part of what turns an initial spark into a fully fledged crisis is that it has not been expected by market participants and regulators. Crises have almost always come as a surprise to most market participants, catching them off-guard and ill-prepared. Goodhart and Persaud (2008) point out that if market prices for assets or credit were good at predicting crashes, they would not happen. Similarly, Abreu and Brunnermeier (2003) explain how asset market bubbles can build up over time despite the presence of rational arbitrageurs. Mispricing can persist in particular during late stages of a bubble. As a result, building early warning signals on price data, or credit risk measurements derived from equity data, has obvious drawbacks.

This section attempts to obtain warning signals for macroprudential policy by investigating the residual dynamics in conventional models of portfolio credit risk. Rather than to focus on what is easily measured, such as the impact of observed macroeconomic and financial risk factors on defaults, we suggest looking for structure in *what is missing*. In particular, we examine the residual, or ‘unexpected’, effects after observed risk factors have been taken into account.

Our warning signals are based on recent research by Koopman, Lucas, and Schwaab (2008) and Duffie, Eckner, Horel, and Saita (2009) who find substantial evidence for a dynamic unobserved risk factor driving default for U.S. firms in addition to observed macroeconomic data, financial market covariates, and firm-specific information such as ratings, equity returns and volatilities. A frailty factor implies that systematic default

risk ('the default cycle') can decouple from what is implied by macroeconomic and financial market conditions ('the business cycle'). In Section 5.4.1 we confirmed that such decoupling is present in both U.S. and non-U.S. data.

Figure 5.8 presents four estimated frailty factors for the U.S. (top left graph), EU-27 countries (top right), Asia (bottom left), and other countries (bottom right). For the U.S., frailty effects have been pronounced during bad times, such as the savings and loan crisis in the U.S. in the late 1980s, leading up to the 1991 recession. They have also been pronounced in exceptionally good times, such as the years 2005-07 leading up to the recent financial crisis. In these years, default conditions are much more benign than would be expected from observed macro and financial data. The in-sample standard error bands suggest that frailty effects are significantly different from zero at these times. As expected from Table 5.4, European frailty effects are negligible. The frailty factors for Asia and the other countries (bottom graphs) are similar negative during the pre-crisis period as in the U.S. This suggests a similar build-up of systematic risk in these areas during this time.

Figure 5.9 presents a decomposition of systematic default conditions into macro, frailty, industry-specific (contagion), and idiosyncratic risk shares based on a rolling window of two years for different economic regions. The decomposition is based on improvements in the McFadden R-squared statistic, see Section 5.2.3. For the U.S., again, frailty effects are pronounced during the savings and loan crisis and leading up to the 1991 recession, and even more so in 2005-07 leading up to the recent crisis. The idiosyncratic risk component decreases over the last ten years, implying higher systematic risk in credit exposures. Such dependence undermines positive effects from diversification at the firm level. It may further lead to a 'fallacy of composition' at the systemic level, as pointed out by e.g. Brunnermeier et al. (2009). If banks react in unison to a shock to highly correlated traded exposures, fire sales and liquidity spirals can ensue, see e.g. Brunnermeier and Pedersen (2009).

The bottom graph of Figure 5.9 presents estimated risk shares for economic regions other than the U.S. Due to the very limited amount of default data in the rolling window in these regions, we focus on the decomposition results for the years 2000Q1 to 2009Q4. For E.U. countries, there is no apparent decrease in the idiosyncratic risk share as in the U.S. Systematic risk increases in the financial crisis (2008Q1 to 2009Q4) mainly due to pronounced changes in observed European macroeconomic data. This is reflected by the share of macro risk. The risk decomposition for the Asian firms in our sample is volatile, but indicates a buildup of systematic risk in the pre-crisis period of 2003-07 also in Asian countries.

Figure 5.8: Frailty factor estimates

The top panel presents conditional mean estimates of four region-specific frailty factors. The estimate is presented with 95% standard error bands. The bottom panel compares smoothed (conditional mode) estimates of frailty effects with the corresponding filtered factors.

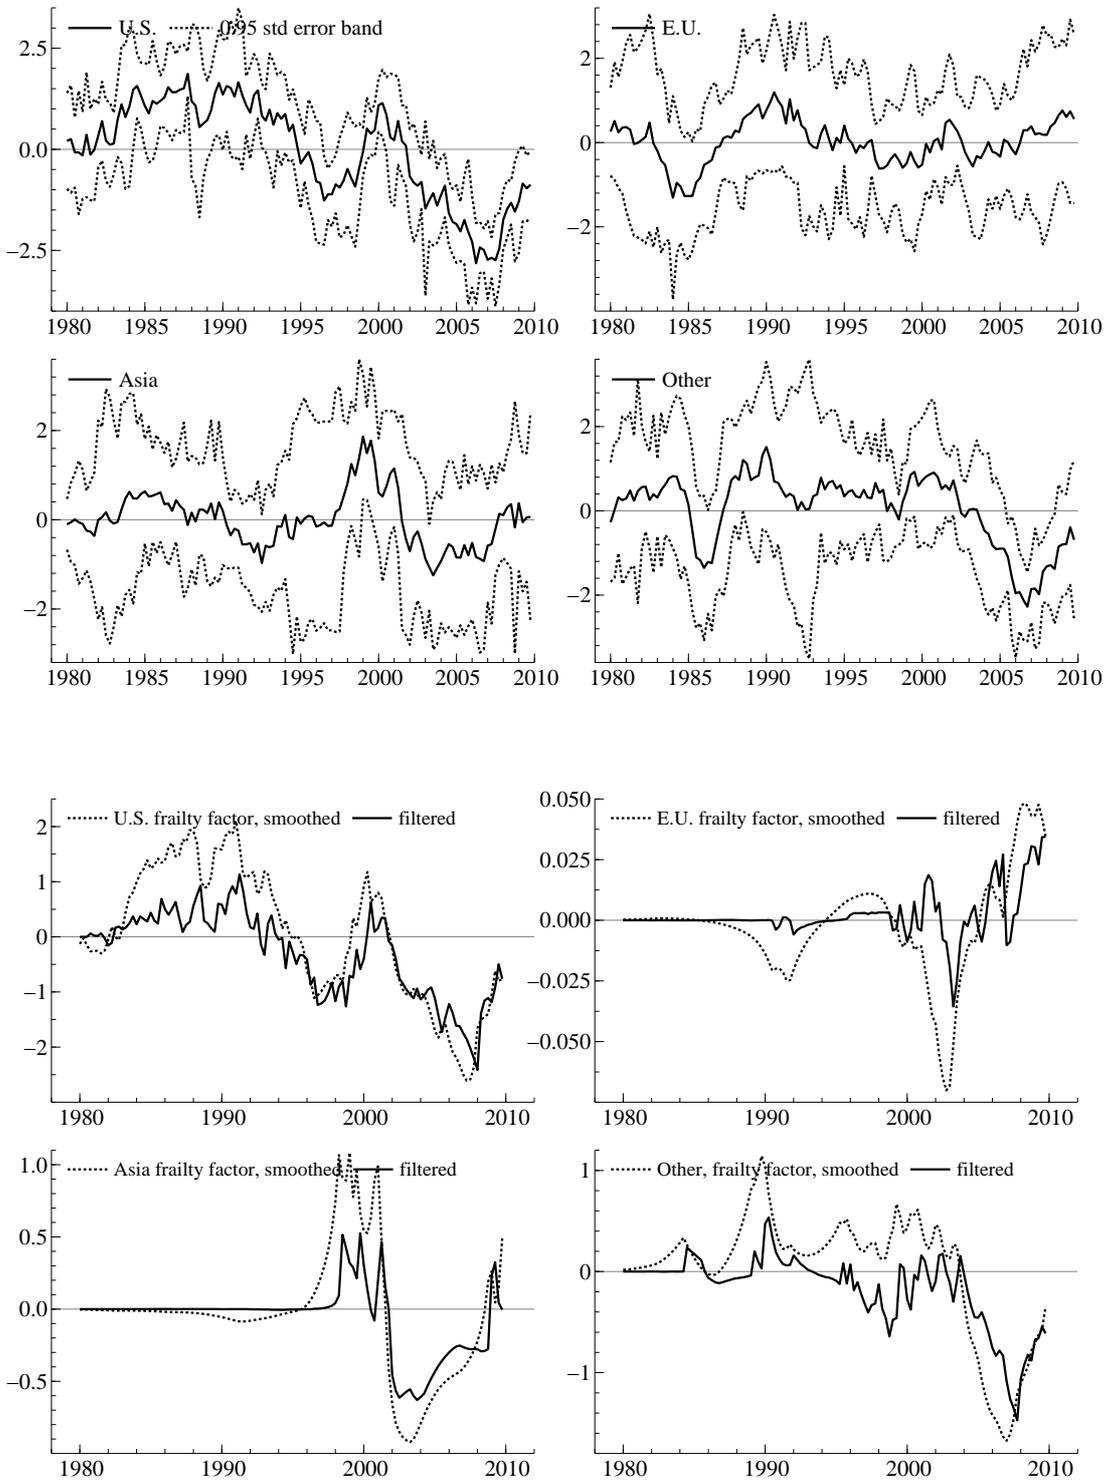
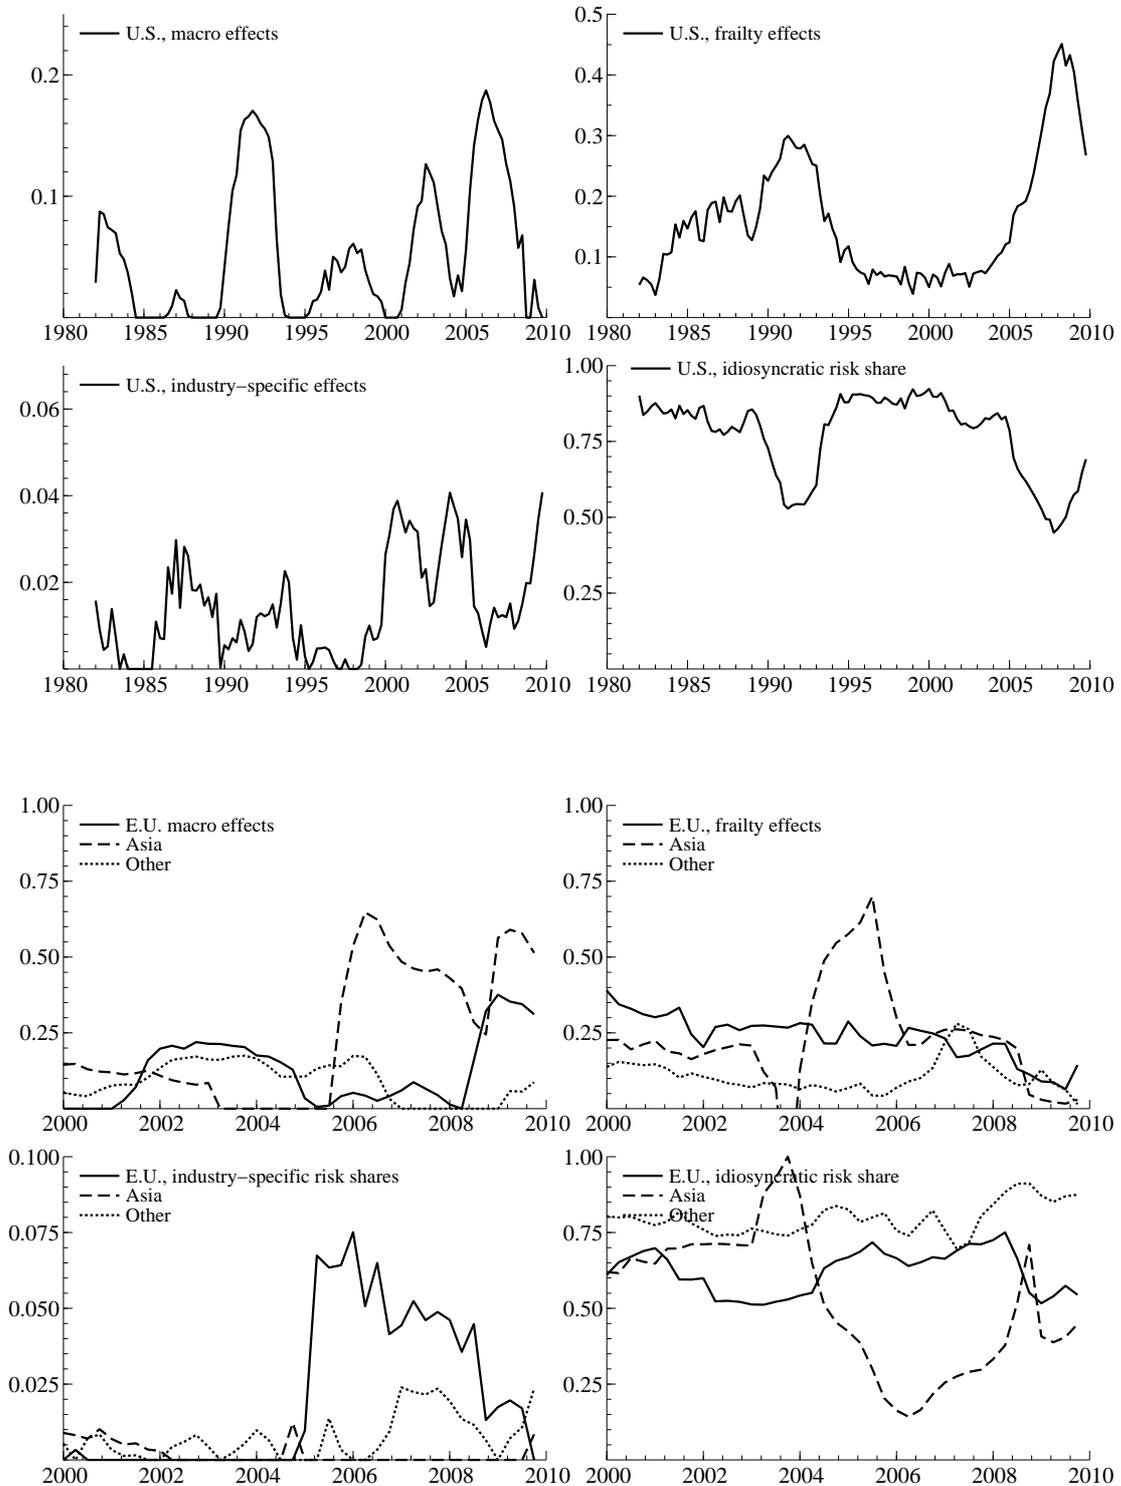


Figure 5.9: Systematic default risk over time

We plot risk shares estimated over a rolling window of eight quarters. The reported sample is from 1980Q1 to 2009Q4 for the U.S. (top graph), and from 2000Q1 to 2009Q4 for E.U., Asian, and other countries (bottom graph).



The finding that systematic credit risk conditions can decouple significantly and persistently from what is implied by macroeconomic and financial market conditions does not imply that such decoupling is always easy to detect in real time. The frailty effects reported in the top graph of Figure 5.8 are based on smoothed risk factor estimates. That means that the complete sample is taken into account for inference on the location of the factors at any time $t = 1, \dots, T$. As a result, smoothed factors use all available information, but look ahead to some extent. The bottom graph of Figure 5.8 compares filtered and corresponding smoothed risk factor estimates. Given parameters, filtered factors are obtained at each point in the sample given past information only. Filtered factors are therefore available in real time. For our data, the filtered estimates track the general pattern of the smoothed factors relatively closely. For example, filtered U.S. frailty effects remain pronounced in the late 80s, and mid-2000s. Nevertheless, they lag the smoothed effects to some extent by construction. Clearly, the out-of-sample assessment of risk conditions in real time is not straightforward, and may require a combination of subjective judgement in addition to model-based predictions.

We conclude that macro-prudential policy makers should be aware that systematic credit risk conditions can decouple significantly and persistently from macroeconomic and financial market conditions. Such decoupling is present also in non-U.S. default data. The timing of frailty effects suggests that such decoupling could be triggered by regime shifts in financial regulation (as preceding e.g. the U.S. savings and loan crisis), or by more gradual innovation in credit risk transfer and securitization in many more countries (as in 2005-07). The monitoring of both business and credit cycle conditions, domestic and abroad, is key if such developments are to be detected.

5.6 Conclusion

We proposed a novel diagnostic framework for financial systemic risk assessment. The model allows us to assess macroeconomic and corporate credit risk conditions in four broad economic regions. In an empirical study of world default data from 1981Q1 to 2009Q4, we estimate latent risk factors and factor sensitivities for firms in the U.S., current EU-27 countries, Asia, as well as other remaining countries. The large-dimensional model allows for the differential impact of international macroeconomic conditions on regional default rates, region-specific frailty factors, as well as industry-specific (contagion) dynamics.

We present an indicator to visualize and compare default stress for a given set of firms. A straightforward measure of financial systemic risk is obtained as a special case. Finally, we suggest that the magnitude of frailty effects at a given point in time can serve

as a warning signal for macro-prudential policy. Aggregate default cycle and business cycle conditions do not need to coincide. Frailty factors capture the divergence of the two processes. Frailty effects have been pronounced during bad times, such as the savings and loan crisis in the U.S. leading up to the 1991 recession, and exceptionally good times, such as the years 2005–07 leading up to the recent financial crisis.

Chapter 6

Conclusion

In this thesis, four articles are presented in four main chapters. Each chapter applies dynamic factor modeling techniques, often based on state space methods, to a credit risk modeling problem at hand.

In Chapter 2, we propose and motivate a novel non-Gaussian panel data time series model with regression effects to estimate and measure the dynamics of corporate default hazard rates. The model is the first to combine a non-Gaussian panel data specification with the principal components of a large number of macroeconomic covariates. The model integrates different types of factors, i.e., common factors from macroeconomic and financial time series, an unobserved latent component for (discrete, count) default data, and ‘observed’ contagion factors at the industry level. At the same time we allow for standard measures such as equity returns, volatilities, and ratings. In an empirical application, we continue to find a large and significant role for a dynamic frailty component even after taking account of more than 80% of the variation from more than 100 macroeconomic and financial covariates, while controlling for contagion at the industry level and equity returns and volatilities. Our findings support earlier research which points out the need for a latent component to prevent a downward bias in the estimation of extreme default losses on portfolios of U.S. corporate debt. Our results indicate that the presence of a latent factor may not be due to a few omitted macroeconomic covariates, but rather appears to capture different omitted effects at different times.

Chapter 3 introduced a new latent dynamic factor model framework, the mixed-measurement dynamic factor model (MM-DFM), for time series observations from different families of distributions and mixed sampling frequencies. Such models are particularly useful for the analysis of credit risk data, but they have applications also beyond that setting. Parameter and latent factor estimates can be obtained by e.g. Monte-Carlo maximum likelihood methods based on importance sampling. We provide two extensions

of that framework. First, we obtain increased computational speed by collapsing observations into a lower-dimensional space such that less observations are passed through the Kalman Filter and Smoother for each evaluation of the log-likelihood. Second, we consider a less complex observation-driven alternative model, the mixed-measurement generalized autoregressive score model (MM-GAS), where the factors are driven by the scaled score of the (local) log-likelihood. Missing values arise due to mixed frequencies and forecasting, and can be accommodated straightforwardly in either the MM-DFM and MM-GAS framework. In an empirical application of the mixed-measurement framework we model the systematic variation in US corporate default counts and recovery rates from 1982Q1 - 2008Q4. We estimate and forecast intertwined default and recovery risk conditions, and demonstrate how to obtain the predictive credit portfolio loss distribution. While the MM-GAS model is simpler and computationally more efficient than the MM-DFM, we do not find that its reduced complexity comes at the cost of diminished out-of-sample (point) prediction accuracy.

Chapter 4 presents 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 the empirical study for U.S. data, we found that approximately a third of default rate volatility at the industry and rating level is systematic. The share of systematic risk caused by macroeconomic and financial activity ranges from about thirty for speculative grade up to sixty percent for investment grade companies. The remaining share of systematic risk is captured by frailty, closely followed by industry factors. 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. This could be detrimental from a financial stability perspective.

Chapter 5 proposes a framework for the measurement of world default conditions. The model allows us to track credit risk conditions around the globe. In an empirical study of world default data from 1981Q1 to 2009Q4, we estimate model parameters and multiple latent risk factors for firms in four geographical regions. The high-dimensional model allows for the differential impact of world business cycle conditions on regional default rates, additional latent geographical risk factors, as well as a region-specific impact of world-wide latent industry sector dynamics. To visualize latent common default stress for a given set of firms in a given region, we present a novel indicator based on the nonlinear

non-Gaussian factor model for disaggregated defaults. The indicator allows us to take account of macroeconomic, frailty, and industry-specific effects simultaneously. Finally, we suggest that the magnitude of ‘frailty’ effects at a given point in time can serve as an early warning signal for macro-prudential policy makers. Frailty effects have been pronounced during bad times, such as the savings and loan crisis in the US leading up to the 1991 recession, and exceptionally good times, such as the years 2005-07 leading up to the recent financial crisis, when defaults were much lower than implied by macro data.

Bibliography

- Abreu, D. and M. K. Brunnermeier (2003). Bubbles and crashes. *Econometrica* 71(1), 173–204.
- Adrian, T. and M. Brunnermeier (2009). Covar. *Federal Reserve Bank of New York Staff Reports* 348.
- Ait-Sahalia, Y., J. Cacho-Diaz, and R. J. A. Laeven (2010). Modeling financial contagion using mutually exciting jump processes. *NBER working paper 15850*, 1–50.
- 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.
- Angelini, P., S. Neri, and F. Panetta (2010). Macroeconomic stabilization policies: Grafting macroprudential tools in a macroeconomic framework. *Bank of Italy Discussion Paper*.
- Angeloni, I. and E. Faia (2009). A tale of two policies: prudential regulation and monetary policy with fragile banks. *mimeo*.
- Azizpour, S. and K. Giesecke (2008). Self-Exciting Defaults: Contagion vs. Frailty. *Stanford University, working paper*.
- Bai, J. (2003). Inferential Theory for Factor Models of Large Dimension. *Econometrica* 71(1), 135–171.
- Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica* 70(1), 191–221.
- Bai, J. and S. Ng (2007). Determining the number of primitive shocks in factor models. *Journal of Business and Economic Statistics* 25(1), 52–60(9).
- 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*.

- Black, F. and J. C. Cox (1976). Valuing corporate securities: Some effects of bond indenture provisions. *The Journal of Finance* 31(2), 351–367.
- Boissay, F. and R. Gropp (2010). Trade Credit Defaults and Liquidity Provision by Firms. *working paper*.
- Bruche, M. and C. Gonzalez-Aguado (2009). Recovery Rates, Default Probabilities and the Credit Cycle. *Journal of Banking and Finance*, forthcoming.
- Brunnermeier, M., A. Crocket, C. Goodhart, A. Persaud, and H. Shin (2009). The fundamental principles of financial regulation. *Geneva Report on the World Economy*, No. 11.
- Brunnermeier, M. K. (2009). Deciphering the liquidity and credit crunch 20072008. *Journal of Economic Perspectives* 23(1), 77–100.
- Brunnermeier, M. K. and L. H. Pedersen (2009). Market liquidity and funding liquidity. *Review of Financial Studies* 22(6), 2201–2238.
- 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.
- Connor, G. and R. A. Korajczyk (1986). Performance Measurement with the arbitrage pricing theory: A new framework for analysis. *Journal of Financial Economics* 15(3), 373–394.
- Connor, G. and R. A. Korajczyk (1988). Risk and Return in equilibrium APT: Application of a new test methodology. *Journal of Financial Economics* 21(2), 255–289.
- Connor, G. and R. A. Korajczyk (1993). A test for the number of factors in an approximate factor model. *Journal of Finance* 48(4), 1263–91.
- Covas, F. and S. Fujita (2009). Time-varying capital requirements in a general equilibrium model of liquidity dependence. *Philadelphia Fed Discussion Papers*.
- Creal, D., S. J. Koopman, and A. Lucas (2008). A General Framework for Observation Driven Time-Varying Parameter Models. *Tinbergen Institute Discussion Paper Series* 108/4.
- Creal, D., B. Schwaab, S. J. Koopman, and A. Lucas (2010). Mixed measurement dynamic factor models. *working paper*, available at www.tinbergen.nl/~bschwaab.
- CreditMetrics (2007). CreditMetrics (TM) - Technical Document, RiskMetrics Group. www.riskmetrics.com/pdf/dnldtechdoc/CMTD1.pdf.

- 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).
- de Jong, P. and N. Shephard (1995). The simulation smoother for time series models. *Biometrika* 82, 339–50.
- de Larosiere, J. (2009). *The high level group on financial supervision in the E.U. - Report*. Bruxelles: The European Commission.
- Doornik, J. A. (2007). *Ox: An Object-Oriented Matrix Language*. Timberlake Consultants Press, London.
- Doz, C., D. Giannone, and L. Reichlin (2006). A Quasi Maximum Likelihood Approach for Large Approximate Dynamic Factor Models. *ECB working paper* 674.
- Duffie, D., A. Eckner, G. Horel, and L. Saita (2009). Frailty Correlated Default. *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.
- Düllmann, K. and M. Trapp (2004). Systematic Risk in Recovery Rates - An Empirical Analysis of U.S. Corporate Credit Exposures. *Deutsche Bundesbank working paper*.
- 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 (2000). Time Series Analysis of Non-Gaussian Observations based on State Space Models from both classical and Bayesian perspectives. *Journal of the Royal Statistical Society* 62(1), 3–56.
- 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.
- ECB (2009). The concept of systemic risk. *European Central Bank Financial Stability Review, December 2009*.
- Exterkate, P., D. van Dijk, C. Heij, and P. J. F. Groenen (2010). Forecasting the Yield Curve in a Data-Rich Environment using the Factor-Augmented Nelson-Siegel Model. *Erasmus University Rotterdam, Econometric Institute working paper*.
- 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*.

- Forni, M., M. Hallin, M. Lippi, and L. Reichlin (2000). The Generalized Dynamic-Factor Model: Identification and Estimation. *The Review of Economics and Statistics* 82(4), 540–554.
- Forni, M., M. Hallin, M. Lippi, and L. Reichlin (2005). The Generalized Dynamic-Factor Model: One-Sided Estimation and Forecasting. *Journal of the American Statistical Association* 100, 830–840.
- 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. (1977). The Dynamic Factor Analysis of Economic Time Series. in Aigner, D and Goldberger, A, eds.: *Latent variables in socio-economic models*, North-Holland.
- 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*.
- Goodhart, C. and A. Persaud (2008). A proposal for how to avoid the next crash. *Financial Times UK*, 31 Jan 2008.
- Gupton, G. and R. Stein (2005). LossCalc V2: Dynamic Prediction of LGD. *Moody's KMV Investors Services*.
- Hartmann, P., S. Straetmans, and C. de Vries (2004). Asset market linkages in crisis periods. *Review of Economics and Statistics* 86(1), 313–326.
- Harvey, A. (1989). *Forecasting, Structural Time Series Models, and the Kalman Filter*. Cambridge University Press.
- Harvey, A. C. (1993). *Time Series Models, 2nd edn*. Hemel Hempstead: Harvester Wheatsheaf.
- Jorion, P. and G. Zhang (2007a). Credit correlations from counterparty risk. *University of California at Irvine, working paper*.
- Jorion, P. and G. Zhang (2007b). Good and bad credit contagion: evidence from credit default swaps. *Journal of Financial Economics* 84(3), 860–883.
- Jungbacker, B. and S. J. Koopman (2007). Monte Carlo estimation for nonlinear non-Gaussian state space models. *Biometrika* 94, 827–839.

- Jungbacker, B. and S. J. Koopman (2008). Likelihood-based Analysis of Dynamic Factor Models. *Tinbergen Institute Discussion Paper*.
- 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 (2008). Forecasting Cross-sections of Frailty-correlated Default. *Tinbergen Institute Discussion Paper 029/04*, 1–38.
- Koopman, S. J., A. Lucas, and B. Schwaab (2010). Macro, Industry and Frailty-Effects in Defaults: The 2008 credit crisis in perspective. *Tinbergen Institute Discussion Paper 2010-004/2*, 1–40.
- Koopman, S. J., N. Shephard, and D. Creal (2009). Testing the assumptions behind importance sampling. *Journal of Econometrics* 149(1), 2–10.
- Koopman, S. J., N. Shephard, and J. A. Doornik (2008). *Statistical Algorithms for Models in State Space Form: Ssfpack 3.0*. Timberlake Consultants Press, London.
- Lando, D. (2003). *Credit Risk Modelling - Theory and Applications*. Princeton University Press.
- 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 bankruptcy announcements. *Journal of Financial Economics* 32, 45–60.
- Lawley, D. N. and A. E. Maxwell (1971). *Factor Analysis as a Statistical Method*. American Elsevier Publishing, New York.
- Lee, K. M. (2010). *Filtering Non-linear State Space Models*. Tinbergen Institute Dissertation: Thela Thesis Amsterdam.
- Ludvigson, S. C. and S. Ng (2007). The empirical risk-return relation: a factor analysis approach. *Journal of Financial Economics* 83, 171–222.

- Massimiliano, M., J. H. Stock, and M. W. Watson (2003). Macroeconomic forecasting in the Euro area: Country specific versus area-wide information . *European Economic Review* 47(1), 1–18.
- 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.
- Merton, R. (1974). On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance* 29(2), 449–470.
- on Financial Regulation, T. S. L. G. (2009). Reforming Capital Requirements for Financial Institutions. *Council on Foreign Relations Working Paper*, 1–8.
- Perotti, E. and J. Suarez (2009). Liquidity charges as a macroprudential tool. *CEPR Policy Insight* 40, November.
- 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.
- Repullo, R. and J. Suarez (2009). The procyclical effects of banking regulation. *CEMFI Discussion Paper*.
- Rösch, D. and H. Scheule (2005). A Multifactor Approach for Systematic Default and Recovery Risk. *The Journal of Fixed Income* 15(2), 63–75.
- Sargent, T. and C. Sims (1977). Business cycle modeling without pretending to have too much a priori economic theory. *Federal Reserve Bank of Minneapolis, working paper No 55*.
- Schuerman, T. (2006). What do we know about Loss-given-default? In D. Shimko (Ed.), *Credit Risk Models and Management*. Risk Books, London, UK.
- Schwaab, B., A. Lucas, and S. J. Koopman (2010). World default conditions, macro-prudential oversight, and warning signals. *mimeo, available at www.tinbergen.nl/~bschwaab*.
- Schweppe, F. (1965). Evaluation of likelihood functions for Gaussian signals. *IEEE Transactions on Information Theory* 11, 61–70.
- Segoviano, M. A. and C. Goodhart (2009). Banking stability measures. *IMF Working Paper*.

- Shephard, N. and M. K. Pitt (1997). Likelihood analysis of non-Gaussian measurement time series. *Biometrika* 84, 653–67.
- Stock, J. and M. Watson (2002a). Forecasting Using Principal Components From a Large Number of Predictors. *Journal of the American Statistical Association* 97(460), 1167–1179.
- Stock, J. and M. Watson (2002b). Macroeconomic Forecasting Using Diffusion Indexes. *Journal of Business and Economic Statistics* 20(2), 147–162.
- Stock, J. and M. Watson (2005). Implications of dynamic factor models for VAR analysis. *NBER working paper*.
- Taleb, N. N. (2007). *The Black Swan: The impact of the highly improbable*. London: Penguin.
- Taylor, J. B. (2008). Monetary policy and the state of the economy. *Testimony before the Committee on Financial Services, U.S. House of Representatives, February 26, 2008*.
- Vassalou, M. and Y. Xing (2004). Default risk in equity premiums. *Journal of Finance* 109(2), 831–868.
- Volcker, P. A., J. A. Frenkel, A. F. Neto, and T. Padoa-Schioppa (2009). *Financial Reform: A framework for financial stability*. Washington D.C.: Group of 30.
- West, M. and J. Harrison (1997). *Bayesian Forecasting and Dynamic Models, 2 edn*. New York: Springer-Verlag.
- White, W. R. (2004). Making macroprudential concerns operational. *Speech at a Financial Stability Symposium organised by the Dutch Central Bank, Amsterdam, October 2004*, 1–4.
- White, W. R. (2006). Procyclicality in the financial system: do we need a new macro-financial stabilisation framework? *BIS Working Papers* 193, 1–20.

Samenvatting (Summary in Dutch)

In dit proefschrift worden vier artikelen in vier hoofdstukken gepresenteerd. Elk hoofdstuk past dynamische-factormodelleringstechnieken, vaak gebaseerd op methoden voor toestand-ruimtemodellen, toe op een kredietrisicoprobleem.

In hoofdstuk 2 introduceer ik een nieuw niet-Gaussiaans panel-data-model met regressie-effecten om wanbetalingsfracties voor bedrijven te schatten. Het model combineert een niet-Gaussiaanse panel-data-specificatie met de principale componenten van een groot aantal macro-economische variabelen. Het model omvat verschillende soorten factoren, te weten, gemeenschappelijke factoren van macro-economische en financiële tijdreeksen, een niet-waargenomen (latente) component voor (discrete) aantallen wanbetalingen, en waargenomen factoren voor besmettingsgevaar op het niveau van de industriële sector. Tegelijkertijd gebruik ik maatstaven voor risico zoals aandelenrendementen, volatiliteit, en kredietwaardigheidsbeoordelingen. In een empirische toepassing vind ik dat een belangrijke rol is weggelegd voor een niet-waargenomen dynamische (frailty of kwetsbaarheid) component, zelfs als rekening wordt gehouden met meer dan 80% van de variatie in meer dan 100 macro-economische en financiële variabelen, besmetting op het niveau van de industriële sector, en rendementen op en volatiliteiten van aandelen. Onze bevindingen ondersteunen eerder onderzoek dat op de noodzaak van een latente component wijst. Deze latente component voorkomt een neerwaartse onzuiverheid in de schatting van de extreme verliezen op portefeuilles van Amerikaanse bedrijfsobligaties en leningen. Onze resultaten geven aan dat de aanwezigheid van een latente factor niet te wijten is aan een aantal weggelaten macro-economische variabelen. In plaats daarvan lijkt de latente factor verschillende weggelaten effecten op verschillende tijdstippen weer te geven.

Hoofdstuk 3 introduceert een nieuw model met latente dynamische factoren (MM-DFM) voor tijdreeksen van verschillende families van verdelingen en met een gemengde observatiefrequentie. Dergelijke modellen zijn bijzonder nuttig voor de analyse van kredietrisico's, maar hebben ook andere toepassingen. Parameters en latente factoren kunnen worden geschat door bijvoorbeeld Monte-Carlo-maximale-aannemelijkheidsmethoden gebaseerd op belanggewogen trekkingsmethoden. Ik introduceer twee uitbreidingen. Ten

eerste is het mogelijk de dimensie van de gegevens te reduceren zodat ze kunnen worden weergegeven in een lagerdimensionale ruimte. Op deze manier hoeven minder waarnemingen te worden doorgegeven aan het Kalman-filter en smoother voor elke evaluatie van de logaritme van de aannemelijkheidsfunctie. Ten tweede introduceren wij een minder complex observatie-aangedreven model (MM-GAS). Hier worden de factoren gedreven door de geschaalde score van de (lokale) logaritme van de aannemelijkheidsfunctie. Ontbrekende waarden ontstaan als gevolg van gemengde frequenties, of door het doen van voorspellingen buiten de steekproef. Dergelijke waarnemingen kunnen zonder moeite worden meegenomen in of het MM-DFM of het MM-GAS framework. In een empirische toepassing van het model met gemengde waarnemingen bestudeer ik de systematische variatie in wanbetalingsrisico's en herstellfracties voor bedrijven in de V.S. van 1982Q1 - 2008Q4. Ik schat en voorspel met elkaar verwoven krediet- en herstellrisico's en demonstreer hoe de voorspelverdeling van het verlies van een kredietportefeuille te verkrijgen is. Hoewel het MM-GAS model eenvoudiger en sneller te berekenen is dan het MM-DFM, tonen mijn bevindingen aan dat de verminderde complexiteit niet ten koste gaat van de nauwkeurigheid van (punt)voorspellingen buiten de steekproef.

Hoofdstuk 4 presenteert een systematische ontleding van het debiteurenrisico op basis van een nieuwe model. Waargenomen wanbetalingsfrequenties worden gemodelleerd samen met macro-economische en financiële indicatoren. De daaruit voortvloeiende gegevensverzameling van continue en discrete variabelen wordt geanalyseerd om de oorzaken van het systematische debiteurenrisico te onderzoeken. Door middel van een dynamische factoranalyse meet ik de bijdrage van macro-, kwetsbaarheids-, en industriespecifieke risicofactoren aan de totale variatie in wanbetalingspercentages. In een empirische studie met gegevens uit de V.S. vindt ik dat ongeveer een derde van de volatiliteit in wanbetalingsfracties systematisch is. Het percentage van de systematische risico's dat wordt verklaard door macro-economische en financiële activiteit varieert van ongeveer dertig procent voor bedrijven met een lage kredietwaardigheid tot zestig procent voor bedrijven met een hoge kredietwaardigheid. Het resterende deel van het systematische risico wordt opgevangen door een kwetsbaarheids-factor, op de voet gevolgd door de industrie factoren. Deze bevindingen suggereren dat een beheerder van kredietrisico voor portefeuilles rekening moet houden met alle drie risicobronnen. Typische modellen die alleen rekening houden met macro-economische variabelen nemen een substantieel deel van het systematische risico niet mee. Dit zou schadelijk kunnen zijn vanuit het oogpunt van financiële stabiliteit.

Hoofdstuk 5 introduceert een model voor het meten van internationale wanbetalingsrisico's. Het model stelt ons in staat kredietrisico's over de hele wereld te volgen. In een empirische studie van wereldwijde kredietgegevens over de periode 1981KW1 -

2009KW4 schatten we modelparameters en meerdere latente risicofactoren voor ondernemingen in vier geografische regio's. Binnen dit hoogdimensionale model modelleer ik hoe de wereldconjunctuurcyclus verschillende invloeden heeft op de regionale wanbetalingsfrequenties. Ook bevat het model aanvullende latente geografische risicofactoren en de regio-specifieke industriële sectordynamiek. Om extreme situaties van gemeenschappelijke wanbetalingsrisico voor een bepaalde categorie van bedrijven in een bepaalde regio te visualiseren, presenteer ik een nieuwe indicator. Deze is gebaseerd op het niet-lineaire niet-Gaussiaans factormodel. De indicator maakt het mogelijk rekening te houden met macro-economische, kwetsbaarheids-, en industrie-specifieke effecten tegelijk. Ten slotte stel ik voor dat de omvang van kwetsbaarheidseffecten op een bepaald punt in de tijd kan dienen als een waarschuwingssignaal voor macro-prudentiële beleidsmakers. Kwetsbaarheidseffecten zijn vooral aanwezig tijdens slechte tijden, zoals de crisis in de V.S. in de aanloop naar de recessie van 1991, en uitzonderlijk goede tijden, zoals de jaren 2005-07 in de aanloop naar de recente financiële crisis, toen het aantal faillissementen lager lag dan geïmpliceerd door macrogegevens.

The Tinbergen Institute is the Institute for Economic Research, which was founded in 1987 by the Faculties of Economics and Econometrics of the Erasmus University Rotterdam, University of Amsterdam and VU University Amsterdam. The Institute is named after the late Professor Jan Tinbergen, Dutch Nobel Prize laureate in economics in 1969. The Tinbergen Institute is located in Amsterdam and Rotterdam. The following books recently appeared in the Tinbergen Institute Research Series:

- 438. R.P. WOLTHOFF, *Essays on Simultaneous Search Equilibrium*.
- 439. Y.-Y. TSENG, *Valuation of travel time reliability in passenger transport*.
- 440. M.C. NON, *Essays on Consumer Search and Interlocking Directorates*.
- 441. M. DE HAAN, *Family Background and Children's Schooling Outcomes*.
- 442. T. ZAVADIL, *Dynamic Econometric Analysis of Insurance Markets with Imperfect Information*.
- 443. I.A. MAZZA, *Essays on endogenous economic policy*.
- 444. R. HAIJEMA, *Solving large structured Markov Decision Problems for perishable-inventory management and traffic control*.
- 445. A.S.K. WONG, *Derivatives in Dynamic Markets*.
- 446. R. SEGERS, *Advances in Monitoring the Economy*.
- 447. F.M. VIEIDER, *Social Influences on Individual Decision Making Processes*.
- 448. L. PAN, *Poverty, Risk and Insurance: Evidence from Ethiopia and Yemen*.
- 449. B. TIEBEN, *The concept of equilibrium in different economic traditions: A Historical Investigation*.
- 450. P. HEEMEIJER, *Expectation Formation in Dynamic Market Experiments*.
- 451. A.S. BOOIJ, *Essays on the Measurement Sensitivity of Risk Aversion and Causal Effects in Education*.
- 452. M.I. LÓPEZ YURDA, *Four Essays on Applied Micro econometrics*.
- 453. S. MEENTS, *The Influence of Sellers and the Intermediary on Buyers' Trust in C2C Electronic Marketplaces*.
- 454. S. VUJIĆ, *Econometric Studies to the Economic and Social Factors of Crime*.
- 455. F. HEUKELOM, *Kahneman and Tversky and the Making of Behavioral Economics*.
- 456. G. BUDAI-BALKE, *Operations Research Models for Scheduling Railway Infrastructure Maintenance*.
- 457. T.R. DANIËLS, *Rationalised Panics: The Consequences of Strategic Uncertainty during Financial Crises*.
- 458. A. VAN DIJK, *Essays on Finite Mixture Models*.
- 459. C.P.B.J. VAN KLAVEREN, *The Intra-household Allocation of Time*.
- 460. O.E. JONKEREN, *Adaptation to Climate Change in Inland Waterway Transport*.
- 461. S.C. GO, *Marine Insurance in the Netherlands 1600-1870, A Comparative Institutional Approach*.

462. J. NIEMCZYK, *Consequences and Detection of Invalid Exogeneity Conditions.*
463. I. BOS, *Incomplete Cartels and Antitrust Policy: Incidence and Detection*
464. M. KRAWCZYK, *Affect and risk in social interactions and individual decision-making.*
465. T.C. LIN, *Three Essays on Empirical Asset Pricing.*
466. J.A. BOLHAAR, *Health Insurance: Selection, Incentives and Search.*
467. T. FARENHORST-YUAN, *Efficient Simulation Algorithms for Optimization of Discrete Event Based on Measure Valued Differentiation.*
468. M.I. OCHEA, *Essays on Nonlinear Evolutionary Game Dynamics.*
469. J.L.W. KIPPERSLUIS, *Understanding Socioeconomic Differences in Health An Economic Approach.*
470. A. AL-IBRAHIM, *Dynamic Delay Management at Railways: A Semi-Markovian Decision Approach.*
471. R.P. FABER, *Prices and Price Setting.*
472. J. HUANG, *Education and Social Capital: Empirical Evidences from Microeconomic Analyses.*
473. J.W. VAN DER STRAATEN, *Essays on Urban Amenities and Location Choice.*
474. K.M. LEE, *Filtering Non Linear State Space Models: Methods and Economic Applications.*
475. M.J. REINDERS, *Managing Consumer Resistance to Innovations.*
476. A. PARAKHONYAK, *Essays on Consumer Search, Dynamic Competition and Regulation.*
477. S. GUPTA, *The Study of Impact of Early Life Conditions on Later Life Events: A Look Across the Individual's Life Course.*
478. J. LIU, *Breaking the Ice between Government and Business: From IT Enabled Control Procedure Redesign to Trusted Relationship Building.*
479. D. RUSINOVA, *Economic Development and Growth in Transition Countries.*
480. H. WU, *Essays on Top Management and corporate behavior.*
481. X. LUI, *Three Essays on Real Estate Finance*
482. E.L.W. JONGEN, *Modelling the Impact of Labour Market Policies in the Netherlands*
483. M.J. SMIT, *Agglomeration and Innovations: Evidence from Dutch Microdata*
484. S.VAN BEKKUM, *What is Wrong With Pricing Errors? Essays on Value Price Divergence*
485. X. HU, *Essays on Auctions*
486. A.A. DUBOVIK, *Economic Dances for Two (and Three)*
487. A.M. LIZYAYEV, *Stochastic Dominance in Portfolio Analysis and Asset Pricing*