

Modeling global financial sector stress and credit market dislocation*

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

^(a) European Central Bank, Financial Research

^(b) Department of Econometrics, VU University Amsterdam

^(c) Department of Finance, VU University Amsterdam

^(d) Tinbergen Institute

May 28, 2013

*Author information: Bernd Schwaab, European Central Bank, Kaiserstrasse 29, 60311 Frankfurt, Germany, bernd.schwaab@ecb.int; Siem Jan Koopman, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, s.j.koopman@feweb.vu.nl; Andre Lucas, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, a.lucas@vu.nl. An earlier version of this paper has been circulated as “Systemic risk diagnostics: coincident indicators and early warning signals”. We would like to thank seminar participants at Bundesbank, Bank of England, University College Dublin, Erasmus University Rotterdam, European Central Bank, the European Systemic Risk Board secretariat, International Monetary Fund, NYU Stern, SUERF conferences in Bruxelles and Helsinki, Tinbergen Institute, VU University Amsterdam, and ZEW Mannheim. We thank in particular conference participants at the joint NY Federal Reserve and NYU Stern “Global systemic risk” conference, the “Extreme dependence in financial markets” conference in Rotterdam, the 2011 “EC-squared” conference in Florence, and the 2011 Annual Cambridge-TI/DSF-UPenn meeting in Amsterdam. We are grateful to four anonymous referees, Tobias Adrian, Dion Bongaerts, Christian Brownlees, Michael Dempster, Robert Engle, Philipp Hartmann, and Michel van der Wel for comments. Andre Lucas thanks the Dutch National Science Foundation (NWO) for financial support. The views expressed in this paper are those of the authors and they do not necessarily reflect the views or policies of the European Central Bank or the European System of Central Banks.

Modeling global financial sector stress and credit market dislocation

Abstract

We introduce a new international model for systematic distress risk of financial institutions from the U.S., the European Union area, and the Asia-Pacific region. The model is a partly nonlinear and non-Gaussian factor model in state space form that we estimate by simulated maximum likelihood. Using the estimated model, we construct measures of global financial sector risk and of credit market dislocation, where credit market dislocation is defined as a significant and persistent decoupling of the credit risk cycle from macro-financial fundamentals in one or more regions. We show that such decoupling has in the past preceded episodes of systemic financial distress. Our new measure provides a risk-based indicator of credit conditions and as such complements earlier quantity-based indicators from the literature. In an extensive comparison with such quantity-based systemic risk indicators, we find that the new indicator behaves competitively to the best quantity-based indicators.

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

JEL classification: *G21, C33*

1 Introduction

In this paper, we introduce an international modeling framework that implies measures of global financial sector stress and new measures of credit market dislocation. The framework is based on a high-dimensional, partly nonlinear and non-Gaussian dynamic factor model in state space form as introduced in Koopman, Lucas, and Schwaab (2012), but here extended to the international context and developed with an explicit focus on financial sector firms. We define financial sector stress in terms of the time varying probability of a joint failure of a large number of currently active financial intermediaries in a given economic region. This definition of financial stress is common in policy and academic work, see for example Segoviano and Goodhart (2009), IMF (2009), and Giesecke and Kim (2011). Credit market dislocation, on the other hand, we define as the situation where credit risk and default experience delinks from the underlying macro fundamentals. Our new modeling framework provides explicit measurements for this quantity.

We provide two main contributions to the literature. First, we provide a tractable framework for tracking and analyzing the drivers of systematic credit risk conditions across different geographical regions and industries. Except for studies like Pesaran, Schuermann, Treutler, and Weiner (2006), such a framework is currently largely missing. Historically, clusters of bank defaults have almost invariably turned out to be very costly in real terms. For example, Reinhart and Rogoff (2009, Chapter 10) conclude that a systemic banking crisis in advanced economies from 1945 to 2006 has on average been followed by a 56% drop in real equity prices, a 36% drop in real estate property prices, a 9% drop in real GDP, a 7% increase in the unemployment rate, a 86% increase in the level of government debt, and 16% drop in sovereign rating score, during three years following such a crisis. In the same chapter, Reinhart and Rogoff also stress that *each* of the ‘big five’ systemic banking crises in advanced economies since World War II (Spain, 1977; Norway, 1987; Finland and Sweden, 1997; and Japan, 1992) has been preceded by a credit boom and an associated increase in asset prices. Both historical observations, i.e., the high cost of crises in real terms as well

as the role of credit markets during the years leading up to the crisis, have been confirmed since then; see the banking crises since 2007 in countries such as Ireland, Spain, and the U.S. As a result, tracking credit market activity over time, which includes credit quantities as well as time-varying systematic credit risk conditions, seems a sensible starting point for financial stability policy analysis at the macro level. A tractable framework for such analysis is indispensable, which is precisely what we develop in this paper.

Second, we analyze new measures of credit market dislocation. These measures are defined as a significant and persistent decoupling of systematic credit risk conditions (the credit cycle) from macro-financial fundamentals (the business cycle). For example, credit risk conditions can be very different from what they are expected to be based on macro-financial covariances during a credit boom. In a credit boom, even bad risks have ample access to credit and can thus avoid or delay default. Consequently, bad risks default less than what would otherwise be expected. Conversely, credit risk conditions may deviate from macro conditions in a credit crunch, when even financially healthy firms find it hard to roll over their debt, thus raising their risk of illiquidity. To our knowledge, the connection between ease of credit access and physical credit risk conditions was first argued informally in Das, Duffie, Kapadia, and Saita (2007), and then more formally in theoretical and empirical models such as for example Duffie, Eckner, Horel, and Saita (2009), He and Xiong (2012) and Koopman, Lucas, and Schwaab (2012). The finding that easily observed macro-financial covariates and firm-specific information, while helpful, are not sufficient to fully explain time-varying systematic credit risk conditions, is a recurring finding in the portfolio credit risk literature, see for example Das, Duffie, Kapadia, and Saita (2007), Koopman, Lucas, and Monteiro (2008), Duffie, Eckner, Horel, and Saita (2009), Koopman, Kräussl, Lucas, and Monteiro (2009), Azizpour, Giesecke, and Schenkler (2010), and Koopman, Lucas, and Schwaab (2011). In this study we document that a decoupling of systematic credit risk conditions from macro fundamentals has in the past preceded macroeconomic and financial distress in both U.S. and non-U.S. data for both financial and non-financial firms. We estimate such decoupling using our state space model and transform it into an indicator. We compare the

performance of this risk-based indicator with a number of more traditional quantity-based measures of credit market conditions and risk and find that the new indicator has a robust and competitive performance to the best quantity-based measures. As such, it provides a complementary perspective to traditional indicators from the literature.

Our work is related to two different streams of literature. First, we relate our empirical study to the growing literature on the construction of financial sector risk measures, see for example Adrian and Brunnermeier (2009), Huang, Zhou, and Zhu (2009), Acharya, Pedersen, Philippon, and Richardson (2010), Brownlees and Engle (2010), Moreno and Pena (2010), and Billio, Getmansky, Lo, and Pelizzon (2012). From a systemic perspective, financial firm defaults are much worse if they occur in clusters. What seems manageable in isolation may not be easily dealt with if the entire system is under stress; see Acharya et al. (2010). Systemic risk is a multi-faceted notion, and can be broadly defined as a set of circumstances that threatens the stability of, or public confidence in, the financial system, as in Billio et al. (2012). A cluster of simultaneous financial firm defaults typically shakes the public confidence in the financial system. Tracking the probability of widespread ‘meltdown’, based on a large cross section of financial firms that constitute the financial system, has therefore a clear role in a prudential financial sector surveillance program.

The construction of financial sector risk measures in the time dimension is extremely challenging. In addition to involving a large cross section of financial firms at any given point in time, also the time-varying marginal risks are unobserved and need to be inferred from available data. For example, Segoviano and Goodhart (2009) adopt a copula perspective to link together the failure of several financial institutions. Their approach is partly non-parametric, whereas our framework is parametric. However, our parametric framework lends itself more easily to extensions to high dimensions, i.e., a large number of individual financial institutions. This is practically impossible in the Segoviano and Goodhart (2009) approach due to the non-parametric characteristics. Extensions to higher dimensions is a relevant issue in our current study, as we take a, literally, global perspective of the financial system. Another paper related to ours is Giesecke and Kim (2011). These authors take a

hazard rate approach with contagion and observed macro-financial factors (no frailty). In contrast to their model setup, our mixed-measurements framework allows us to model the macro developments and default dynamics in a joint factor structure. Giesecke and Kim, by contrast, take the macro data as exogenous regressors in their analysis. Also, our study explicitly incorporates the global dimension, and looks beyond coincident indicators of risk.

Second, we relate to the work on accurately measuring point-in-time credit risk conditions. In general, this is a complicated task since not all processes that determine corporate default and financial sector stress are easily observed. Recent research indicates that readily available macro-financial variables and firm-level information may not be sufficient to capture the large degree of default clustering present in corporate default data, see e.g. Das et al. (2007). In particular, there is substantial evidence for an additional dynamic unobserved ‘frailty’ risk factor as well as contagion dynamics, see McNeil and Wendin (2007), Koopman et al. (2008), Koopman and Lucas (2008), Lando and Nielsen (2009), and Duffie et al. (2009), and Azizpour et al. (2010). ‘Frailty’ and contagion risk cause default dependence above and beyond what is implied by observed covariates alone. Compared to these earlier papers, our current paper takes an explicit international perspective.

To the best of our knowledge, we are the first to investigate the magnitude of frailty effects in non-U.S. credit risk data, and we are only aware of the influential study by Pesaran, Schuermann, Treutler, and Weiner (2006) that takes a similar international perspective on global time-varying macro and credit risk conditions. Our study is different in that our setup distinguishes between macro, frailty, and industry effects in a state space setup. In addition, our focus is on financial sector risk instead of stress testing globally active lending institutions. Finally, our setup provides a unified framework to integrate risk information from different sources, whether macroeconomic and financial market conditions, equity markets and balance sheet information (via expected default frequencies, EDFs), or actual defaults. Particularly EDF data is very helpful in the current context, where financial sector default counts are extremely low and result in substantial model risk. EDF data can help to amplify the signal on systematic financial sector risk.

The remainder of this paper is organized as follows. Section 2 introduces the modeling framework and financial sector risk measures. Section 3 describes the data. Section 4 discusses the empirical results. Section 5 concludes. The Appendix contains some details on the estimation procedure and diagnostic checks. A further web appendix is available with additional methodological details and unreported further empirical results.

2 The modeling framework

2.1 Mixed-measurement dynamic factor models

We consider the following tri-part data structure

$$x_t = (x_{1,1,t}, \dots, x_{1,N_1,t}, \dots, x_{R,1,t}, \dots, x_{1,N_R,t})', \quad (1)$$

$$y_t = (y_{1,1,t}, \dots, y_{1,J_1,t}, \dots, y_{R,1,t}, \dots, y_{R,J_R,t})', \quad (2)$$

$$z_t = (z_{1,1,t}, \dots, z_{1,S_1,t}, \dots, z_{R,1,t}, \dots, z_{R,S_R,t})', \quad (3)$$

where $x_{r,n,t}$ represents the n th, $n = 1, \dots, N_r$, macroeconomic or financial markets variable for region $r = 1, \dots, R$ measured at time $t = 1, \dots, T$; $y_{r,j,t}$ is the number of defaults between times t and $t + 1$ for economic region r and cross section $j = 1, \dots, J_r$; and $z_{r,s,t}$ is the expected default frequency (EDF) of financial firm $s = 1, \dots, S_r$ in economic region r at time t . Cross section j can represent different categories of firms, such as industry sectors, rating categories, firm age cohorts, or a combination of these. The model thus includes more ‘standard’, possibly normally distributed macro and financial markets variables x_t , but also count variables y_t and variables z_t that are bounded to the $[0,1]$ interval. Also note that the panel (x_t, y_t, z_t) is typically unbalanced, such that variables may not be observed at all times.

Rather than only using historical firm default data y_t for measuring financial sector credit risk, we also use EDF data z_t for financial firms in our current study. This makes up for a clear lack of historical default experience for such firms in typical data sets such as those of Moody’s or S&P. For example in our empirical sample later on, we count only 12 defaults for European

financial firms. EDFs are proprietary estimates of physical time-varying default probabilities and are based on structural models for credit risk that integrate information from accounting data (via debt levels) and equity markets (via prices and volatilities). Therefore, they are less prone to the criticism that the estimates also embed a strong risk premium component, as would for example be the case if CDS spread data were used instead. EDF data, along with associated distance-to-default measures, are standard conditioning variables in the credit risk literature, see for example Duffie, Saita, and Wang (2007), Bharath and Shumway (2008), and Duffie et al. (2009).

To integrate the different types of variables x_t , y_t , and z_t in one unifying framework, we assume that all variables are driven by a vector f_t of common dynamic factors. We adopt a standard conditional independence assumption: conditional on the common risk factors f_t , the measurements (x_t, y_t, z_t) are independent over time and in the cross section. We make the following modeling assumptions:

$$x_{r,n,t}|f_t^m \sim \text{Gaussian}(\tilde{\mu}_{r,n,t}, \tilde{\sigma}_{r,n}^2), \quad (4)$$

$$y_{r,j,t}|f_t^m, f_t^d, f_t^i \sim \text{Binomial}(k_{r,j,t}, \pi_{r,j,t}), \quad (5)$$

$$\bar{z}_{r,s,t}|f_t^m, f_t^d, f_t^i \sim \text{Gaussian}(\bar{\mu}_{r,s,t}, \bar{\sigma}_{r,s}^2), \quad \bar{z}_{r,s,t} = \ln(z_{r,s,t}/(1 - z_{r,s,t})), \quad (6)$$

with

$$\pi_{r,j,t} = (1 + e^{-\theta_{r,j,t}})^{-1}, \quad (7)$$

$$\tilde{\mu}_{r,n,t} = \tilde{c}_{r,n} + \tilde{\beta}'_{r,n} f_t^m, \quad (8)$$

$$\theta_{r,j,t} = \lambda_{r,j} + \beta'_{r,j} f_t^m + \gamma'_{r,j} f_t^d + \delta'_{r,j} f_t^i, \quad (9)$$

$$\bar{\mu}_{r,s,t} = \bar{c}_{r,s} + \bar{\beta}'_{r,s} f_t^m + \bar{\gamma}'_{r,s} f_t^d + \bar{\delta}'_{r,s} f_t^i, \quad (10)$$

where the common risk factor $f_t' = (f_t^{m'}, f_t^{d'}, f_t^{i'})$ is partitioned in a component capturing macro-financial risks f_t^m , default specific (frailty) risk f_t^d , and industry risk f_t^i . The industry factors f_t^i may arise as a result of default dependence through up- and downstream business links, and may capture the industry-specific propagation of aggregate shocks. The parameters $\tilde{c}_{r,n}$, $\tilde{\beta}_{r,n}$, $\tilde{\sigma}_{r,n}^2$, $\lambda_{r,j}$, $\beta_{r,j}$, $\gamma_{r,j}$, $\delta_{r,j}$, $\bar{c}_{r,s}$, $\bar{\beta}_{r,s}$, $\bar{\gamma}_{r,j}$, $\bar{\delta}_{r,j}$, and $\bar{\sigma}_{r,s}^2$ are fixed parameters that

need to be estimated.

Though the model appears rather complex due to the combination of continuous and discrete data under a unifying factor structure, the model's components are straightforward to interpret. The macro-financial data x_{nt} follow a conditional (on f_t) standard Gaussian distribution where the means depend linearly on the macro risk factors f_t^m . The default counts $y_{r,j,t}$, by contrast, follow a conditional binomial distribution with probability of default $\pi_{r,j,t}$ and number of exposures $k_{r,j,t}$, where the latter are observed directly from the data. The probabilities of default $\pi_{r,j,t}$ are standard logistic transformations of a baseline level $\lambda_{r,j}$ and linear combinations of both the macro-financial (f_t^m), default specific (f_t^d), and industry (f_t^i) risk factors. Finally, the log-odds transformed EDF data $\bar{z}_{r,s,t}$ follow a conditional normal distribution with the mean depending on all components of f_t . The default specific and industry factors thus absorb any time-variation in default rates and EDFs above and beyond what is already captured by the macro-financial factors f_t^m .

The model is completed by postulating a time-series model for the risk factors f_t . We assume that the elements of f_t have independent autoregressive dynamics,

$$f_t = \Phi f_{t-1} + \eta_t, \quad \eta_t \sim \text{NID}(0, \Sigma_\eta), \quad (11)$$

where the coefficient matrix Φ and covariance matrix Σ_η are diagonal, and $f_t \sim \text{N}(0, \Sigma_0)$. Extensions to more complex dynamic structures are straightforward. The autoregressive structure in (11), however, already allows for sufficient stickiness in the components of f_t given the data at hand. For example, it allows the macroeconomic factors f_t^m to evolve slowly over time and to capture shared business cycle dynamics in macro and default data. Similarly, the credit climate and industry default conditions are modeled as persistent processes f_t^d and f_t^i , respectively. The disturbance vector η_t is serially uncorrelated. To ensure the identification of the factor loadings $\tilde{\beta}_{r,n}$, $\beta_{r,j}$, $\bar{\beta}_{r,s}$, $\gamma_{r,j}$, $\bar{\gamma}_{r,s}$, $\delta_{r,j}$, and $\bar{\delta}_{r,s}$, we impose $\Sigma_\eta = I - \Phi\Phi'$. It implies that $\text{E}[f_t] = 0$, $\text{Var}[f_t] = I$, and $\text{Cov}[f_t, f_{t-h}] = \Phi^h$, for $h = 1, 2, \dots$. As a result, the loading coefficients in (9) can be interpreted as risk factor volatilities (standard deviations) for the firms in cross section j for region r .

The above model (4)–(11) extends the factor model structure for partly nonlinear and non-Gaussian panel data as introduced in Koopman, Lucas, and Schwaab (2011, 2012) to the international setting. This allows us to study systematic credit risk dynamics and their inter-regional linkages. In particular, the model can be used to obtain estimates of the unobserved risk factors f_t , which can then be used to construct coincident risk indicators and early warning signals. This is explained in more detail in the next subsections.

As the model is rich in parameters, we impose pooling and normalization restrictions to obtain a more parsimonious representation, while still allowing for sufficient flexibility to describe the data well. First, we link the log odds ratios $\theta_{r,j,t}$ and $\bar{\mu}_{r,s,t}$ such that $\theta_{r,j,t} - \lambda_{r,j}$ and $\bar{\mu}_{r,s,t} - \bar{c}_{r,s}$ are proportional. This makes economic sense, as (9) and (10) relate the log-odds of the (discrete) default counts and of the (continuous) EDF data for the same industry-region combination to the same underlying financial distress factors f_t^d and f_t^i . Second, we normalize the $x_{r,n,t}$ and $\bar{z}_{r,s,t}$ variables to have zero mean and unit variance, such that we can impose $\tilde{c}_{r,n} = \bar{c}_{r,s} = 0$. This further reduces the overall number of parameters that need to be estimated. Implied conditional variances for the Gaussian data are $\tilde{\sigma}_{r,n}^2 = 1 - \tilde{\beta}'_{r,n} \tilde{\beta}_{r,n}$ and $\bar{\sigma}_{r,n}^2 = 1 - \bar{\beta}'_{r,s} \bar{\beta}_{r,s} - \bar{\gamma}'_{r,s} \bar{\gamma}_{r,s} - \bar{\delta}'_{r,s} \bar{\delta}_{r,s}$, and take values on the unit interval. Finally, we follow Koopman and Lucas (2008) and impose an additive structure on the parameters that are cross section and region pair specific, such as

$$\beta_{r,j} = \beta_0^c + \beta_{1,d_j}^c + \beta_{2,r_j}^c + \beta_{3,s_j}^c, \quad (12)$$

where β_0^c represents the baseline effect, $\beta_{1,d}^c$ is the industry-specific deviation, $\beta_{2,r}^c$ is a deviation related to regional effects, and $\beta_{3,s}^c$ is a deviation related to rating group. Similar additive structures are assumed for other coefficients. As usual, some of the dummy effects in (12) have to be omitted for identification. We refer to Section 4.1 for more details.

2.2 Parameter and risk factor estimation by Monte Carlo maximum likelihood

The mixed measurement dynamic factor model presented in the previous section is an extension of the non-Gaussian measurement state-space models as discussed in Shephard and Pitt (1997) and Durbin and Koopman (1997) to modeling observations from different families of parametric distributions. A local version of the model for U.S. default counts was estimated in Koopman, Lucas, and Schwaab (2012). The current inter-regional version that also includes the EDF data is more extensive to estimate, but uses a similar estimation set-up as for the local model. We therefore keep the present methodology discussion short and non-technical, referring to Koopman et al. (2012) and the Appendix for further details. The web appendix to this paper gives details on the treatment of missing values in the current data set.

The model relies on a parameter vector ψ that contains the coefficients in Φ , $\tilde{\beta}_{r,n}$, $\lambda_{r,j}$, $\beta_{r,j}$, $\gamma_{r,j}$, $\delta_{r,j}$, $\bar{\beta}_{r,s}$, $\bar{\gamma}_{r,s}$, and $\bar{\delta}_{r,s}$. This parameter vector is estimated by the method of simulated maximum likelihood. Since the model contains unobserved components f_t and our dynamic factor model is partly non-Gaussian, the likelihood function is not available in closed form. Instead, we have to integrate out all unobserved components from the joint likelihood for ψ and f_t . To evaluate this high-dimensional integral, numerical integration is computationally infeasible. Therefore, we rely on Monte Carlo simulation methods for evaluating the likelihood function. We refer to Koopman et al. (2012) and the Appendix for details on our simulation based estimation procedure for mixed measurement panel data.

2.3 Measures of financial sector stress

Using the mixed measurement model set-up of Section 2.1, we can easily construct estimates of financial sector stress for a specific region or combination of regions based on the model's parameter estimates and estimated risk factors f_t . Such measurements automatically integrate the effects of macro, frailty, and industry effects as represented in f_t . Systematic risk

due to shared exposure to the common risk factors f_t cannot be diversified in a large cross section. Therefore, these factors constitute the main drivers of the time-varying risk of joint financial firm defaults.

Our coincident measures of risk all summarize the information on the time-varying (conditional) probability of default $\pi_{r,j,t}$. The most straightforward indicator is the model-implied financial sector expected failure rate $\text{EFR}_{r,j,t} = \pi_{r,j,t}$ itself, which can be interpreted as the fraction of firms of type j at time t that are expected to fail over the next three months. High failure rates imply high levels of common financial distress, and thus a higher risk of adverse effects on the real economy through financial firm defaults. The measure can be aggregated across sectors or even regions using for example an exposure-weighted average, such as

$$\overline{\text{EFR}}_{r,j',t} = \left(\sum_{j \in j', r \in r'} k_{r,j,t} \right)^{-1} \sum_{j \in j', r \in r'} k_{r,j,t} \pi_{r,j,t}, \quad (13)$$

where j' and r' collect the cross sections and regions of interest.

A second straightforward measure that is directly derived from $\pi_{r,j,t}$ is the joint probability of default (JPoD). The joint probability of default can easily be constructed from the binomial cumulative distribution function and the time-varying financial sector failure rates. For example, the joint probability of default of at least n^* firms of type j at time t is given by

$$\text{JPoD}_{r,j,t}(n^*) = 1 - \sum_{n=0}^{n^*-1} \text{Binomial}(n; k_{r,j,t}, \pi_{r,j,t}), \quad (14)$$

where the $\text{Binomial}(\cdot)$ is the binomial probability mass function. Naturally, a high joint probability of default indicates adverse financial conditions.

A final coincident indicator is the probability integral transform of the log-odds ratio $\theta_{r,j,t}$ in (9), where $\theta_{r,j,t} = \log(\pi_{r,j,t}) - \log(1 - \pi_{r,j,t})$ and where $\pi_{r,j,t}$ is a probability of default. Equation (9) implies that $\theta_{r,j,t}$ is a linear combination of normally distributed random variables, and thus itself normally distributed. The probability integral transform is

$$\text{SSS}_{r,j,t} = \text{N}((\theta_{r,j,t} - \text{E}(\theta_{r,j,t}))/\text{Var}(\theta_{r,j,t})^{1/2}), \quad (15)$$

where $E(\theta_{r,j,t}) = \lambda_{r,j}$, is the unconditional mean of the log-odds ratio, $\text{Var}(\theta_{r,j,t}) = \beta'_{r,j}\beta_{r,j} + \gamma'_{r,j}\gamma_{r,j} + \delta'_{r,j}\delta_{r,j} \geq 0$ is the unconditional variance of $\theta_{r,j,t}$, and $N(\cdot)$ is the standard normal cdf. Values of $\text{SSS}_{r,j,t}$ lie between 0 and 1 by construction with uniform unconditional probability. Values below 0.5 indicate less-than-average (over time) common default stress, while values above 0.5 indicate above-average stress. Values below 20%, say, are historically exceptionally benign, and values above 80% are indicative of substantial systematic stress. Such a historical comparison of stress levels is useful because different financial stability concerns are relevant in different economic environments. Policy makers want to mitigate the buildup of financial imbalances such as asset and credit bubbles; this is particularly a concern during good times, see Borio (2010). Conversely, a credit crunch due to financial sector deleveraging is a concern in particular when financial sector stress is high. Our measure of financial stress is obtained when (15) is applied to model-implied failure rates for financial firms in a given region or block of regions.

2.4 Credit risk deviations

Severe financial bubbles, such as the dot-com asset bubble from approximately 1997-2000 and the credit and real estate bubble in the United States from approximately 2005-2007, occurred when perceptions regarding financial sector risk were below average and even unusually low, see for example Aoki and Nikolov (2012) and Boissay, Collard, and Smets (2013). This is in line with the “financial stability paradox” of Borio (2010, 2011) and the “volatility paradox” of Brunnermeier and Sannikov (2013), which state that the financial sector looks strongest precisely when risk taking is highest and the sector as a whole is the most vulnerable. Credit growth and asset prices are unusually strong, leverage measured at market prices is artificially low, and also risk premia and volatilities are unusually low precisely when financial sector risk is highest. What looks like low risk is, in fact, a sign of aggressive risk-taking by market participants. If the financial stability paradox applies, then an early warning indicator must, perhaps counter-intuitively, not flag model-implied risk, but flag the apparent absence of it and quantify deviations from where risk ‘should be’

based on fundamentals. This section presents such a risk deviation measure.

The coincident risk indicators introduced in the previous subsection take account of all credit risk factors in f_t . However, the structure of our model allows us to make a distinction between the macro factors f_t^m and the other factors (f_t^d, f_t^i) . In particular, the factors f_t^d and f_t^i provide a signal whether local default experience in a particular industry and region is unexpectedly different from what would be expected based on macro fundamentals f_t^m . As such, the two factors can be combined to measure a decoupling of credit risk conditions from their macro-financial fundamentals. This can easily be done by constructing a ‘credit risk deviations’ (CRD) early warning indicator, defined as

$$\text{CRD}_{r,j,t} = (\gamma'_{r,j} f_t^d + \delta'_{r,j} f_t^i) / \sqrt{\gamma'_{r,j} \gamma_{r,j} + \delta'_{r,j} \delta_{r,j}}. \quad (16)$$

This indicator is standard normally distributed, such that we can easily assess whether it is unexpectedly high or low during a specific time period. Section 4 reports and discusses the CRD indicator in more detail for the data at hand. In particular, we demonstrate that major deviations of credit risk conditions from what is implied by standard macro-financial fundamentals have in the past preceded financial and macroeconomic distress.

3 Default risk data and macro-financial covariates

We use quarterly data from three main sources. First, a panel $x_{r,n,t}$ of macroeconomic and financial time series data is taken from Datastream with the aim to capture international business cycle and financial market conditions. We pick macro-financial covariates that are usually stressed in a macro stress test, see for example CEBS (2010) and Tarullo (2010). The macro panel consists of business cycle indicators (real GDP, industrial production, unemployment rate, and an industrial confidence indicator such as the ISM purchasing managers index), price data (inflation, stock market returns), short term and long term interest rates (implicitly the term structure), as well as residential property prices. These variables are included in the macro panel for both the U.S. and Europe, which yields $2 \times 9 = 18$ variables. To proxy business cycle conditions in the Asia-Pacific region, we include real GDP data for

Japan, South Korea, and Australia. This yields 21 macroeconomic variables in total. We refer to the web appendix for a detailed listing and time series plots of the macro data. The covariates enter the analysis as quarterly growth rates from 1980Q1 to 2012Q2, and are available at a quarterly frequency.

The second dataset is constructed from default and exposure data from Moody's. The database contains rating transition histories and default dates for all rated firms (worldwide) from 1980Q1 to 2012Q2. From these data, we construct quarterly values for $y_{r,j,t}$ and $k_{r,j,t}$ in (5). We distinguish default data from three economic regions. We first collect firm and default data from firms headquartered in the United States. The E.U. region consists of the 27 current member states of the European Union plus Switzerland. The Asia-Pacific region is defined as the respective rest of the world, excluding Africa and the Americas. Figure 1 plots aggregate default counts, exposures, and observed fractions over time for each economic region. When counting exposures $k_{r,j,t}$ and corresponding defaults $y_{r,j,t}$, a previous rating withdrawal is ignored if it is followed by a later default. In this way, we limit the impact of strategic rating withdrawals. If there are multiple defaults per firm, we consider only the first event. We use Moody's broad industry classification to distinguish between financial and non-financial firms. Table 1 provides an overview over the default and exposure counts. Data is most abundant for the U.S., with European countries second. Firms are classified into the financial or non-financial sector, and rated either investment grade (IG) or speculative grade (SG).

For financial firms, we add data from a third dataset on expected default frequencies for 55 large (based on 2008Q4 market cap) financial firms in the United States, European Union area, and Asia-Pacific region. The data are taken from Moody's KMV CreditEdge. The $20 + 20 + 15 = 55$ expected default frequency series $z_{r,s,t}$ are based on a firm value model that takes equity values and balance sheet information as inputs. We use the $z_{r,s,t}$ data to augment our relatively sparse data on actual defaults $y_{r,j,t}$ for financial firms. Table 2 lists the identities of the financial firms whose EDF time series data is included in the default risk panel for our empirical analysis. Figure 2 plots EDF panel data for large global

Figure 1: Actual default experience

We present time series plots of (a) the total default counts $\sum_j y_{r,j,t}$ aggregated to a univariate series, (b) the total number of firms $\sum_j k_{r,j,t}$ in the database, and (c) aggregate default fractions $\sum_j y_{r,j,t} / \sum_j k_{r,j,t}$ over time. We distinguish different economic regions: the United States, the European Union area, and the Asia-Pacific region. Data is from 1980Q1 to 2012Q2.

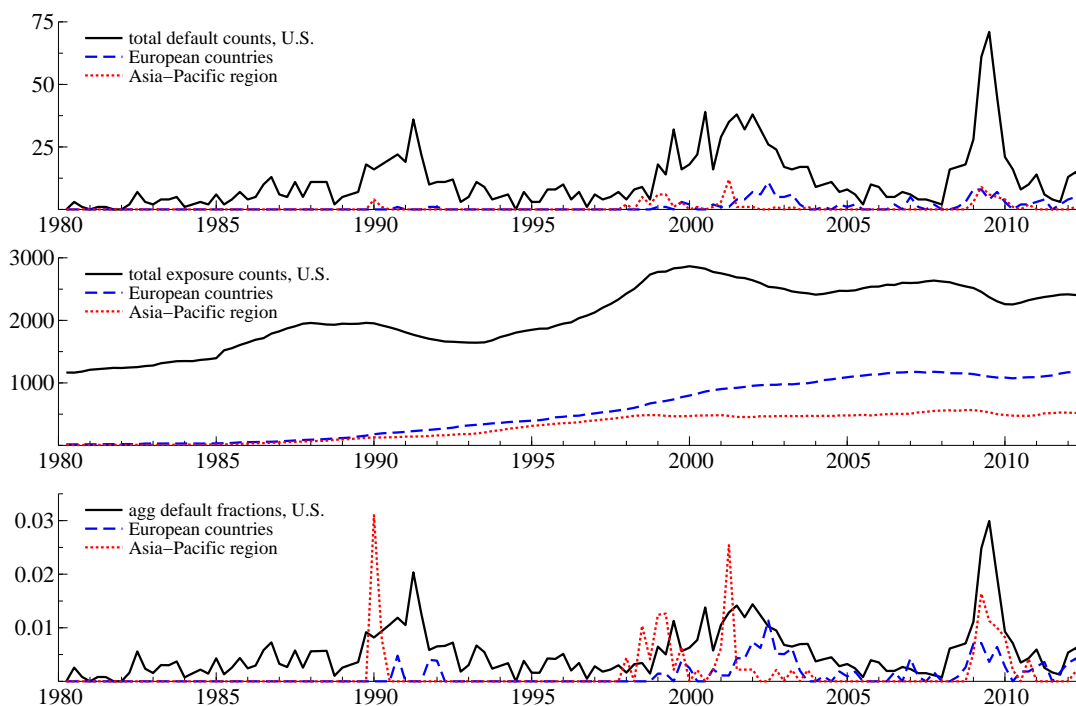


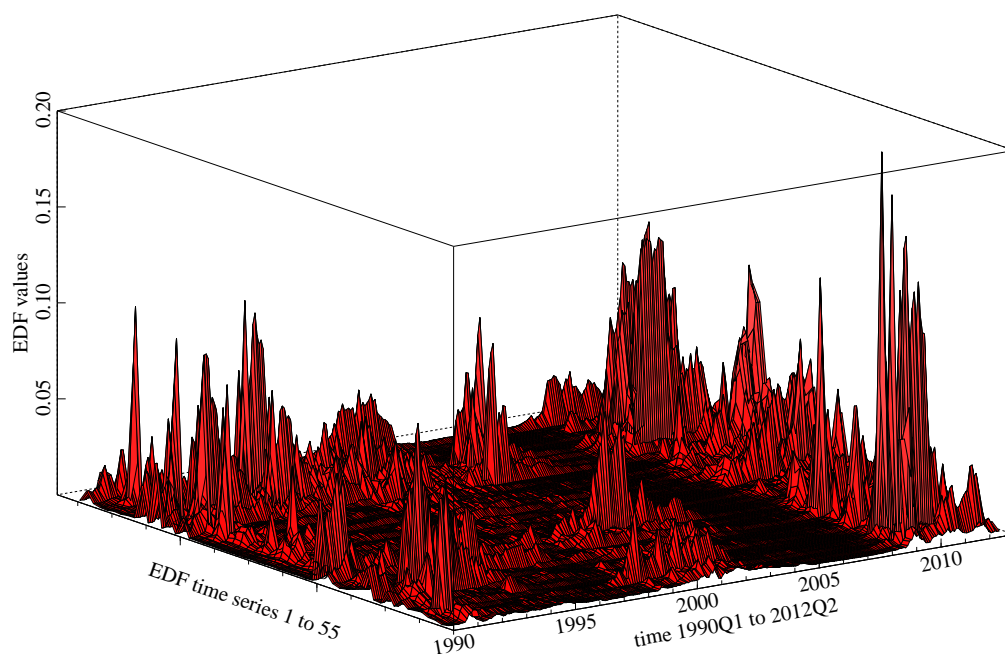
Table 1: International default and exposure counts

The table reports total default counts $\sum_t y_{r,j,t}$ that are disaggregated across economic region, industry sector, and rating group. Exposure counts $k_{r,j,t}$ are reported for selected points in time at ten year intervals.

Region	Industry sector	Rating group	Defaults 1980Q1 to 2012Q2	Exposures 1982Q2	Exposures 1992Q2	Exposures 2002Q2	Exposures 2012Q2
U.S.	fin	IG	9	102	230	439	321
		SG	102	11	66	81	89
	non-fin	IG	20	828	819	1009	731
		SG	1375	305	546	1068	1264
E.U.	fin	IG	8	7	136	359	325
		SG	12	0	3	8	62
	non-fin	IG	2	14	127	428	480
		SG	115	0	6	168	304
A.P.	fin	IG	1	1	87	117	159
		SG	15	0	3	34	51
	non-fin	IG	0	8	71	250	245
		SG	68	0	0	66	65

Figure 2: Expected Default Frequencies (EDFs) of large financial institutions

We plot EDF panel data for 55 large global financial firms from Moody's KMV. The sample consists of the largest 20 U.S. (series 1–20), EU-27 (series 21–40), and rest of the world (series 41–55) financial firms, respectively. The raw data sample is from 1990Q1 to 2012Q2.



financial firms. In the graph we inferred missing values using the EM algorithm of Stock and Watson (2002) to obtain a balanced panel. The figure clearly features the different periods of financial distress, both around the early 1990s recession and the burst of the dotcom bubble in the early 2000s, but particularly during the recent financial crisis.

4 Empirical results

4.1 Model specification

Since we focus almost exclusively in our current paper on financial sector risk, we pool the parameters over all non-financial sectors. Firm-specific ratings are either investment grade or speculative grade. Risk factor sensitivities, however, are allowed to be different across geographical regions and between the financial and non-financial sector.

Table 2: EDF listing

Financial firms whose EDF time series data is included in the default risk panel for our empirical analysis.

U.S.	Europe	Asia-Pacific
AFLAC INC	ALLIANZ SE	AUSTRALIA AND NEW ZL BANKING GRP
ALLSTATE CORP	ASSICURAZIONI GENERALI	BANK OF CHINA LTD
AMERICAN EXPRESS CORP	AXA	CHEUNG KONG (HOLDINGS) LTD
AMERICAN INTERNATIONAL GRP	BBVA SA	CHINA CONSTRUCTION BANK
BANK OF AMERICA CORP	BANCO SANTANDER SA	COMMONWEALTH BANK OF AUSTRALIA
BANK OF NEW YORK MELLON	BARCLAYS PLC	HANG SENG BANK LTD
CITIGROUP INC	BNP PARIBAS	HSBC HOLDINGS PLC
GOLDMAN SACHS GROUP INC	CREDIT AGRICOLE SA	IND & COMRCL BANK OF CHINA
JPMORGAN CHASE & CO	CREDIT SUISSE GROUP AG	MITSUBISHI UFJ FINANCIAL GRP
LOEWS CORP	DEUTSCHE BANK AG	MIZUHO FINANCIAL GRP
MASTERCARD INC	HSBC HOLDINGS PLC	NATIONAL AUSTRALIA BANK LTD
METLIFE INC	ING GROEP N.V.	SBERBANK ROSSEII
MORGAN STANLEY	INTESA SANPAOLO SPA	SUMITOMO MITSUI FINANCIAL
PRUDENTIAL FINANCIAL INC	KBC GROUP NV	TOKIO MARINE HOLDINGS INC
SCHWAB (CHARLES) CORP	ROYAL BANK OF SCOTLAND	WESTPAC BANKING CORP
STATE STREET CORP	SOCIETE GENERALE	
TRAVELERS COS INC	STANDARD CHARTERED PLC	
U S BANCORP	UBS AG	
VISA INC	UNICREDIT SPA	
WELLS FARGO & CO	ZURICH FINANCIAL SERVICES	

For the selection of the factor structure we rely on likelihood-based information criteria (IC). Based on standard criteria such as likelihood, AIC, BIC, and the panel information criteria of Bai and Ng (2002), two or three macro factors are appropriate to summarize the information content in the macro panel. We include three macro factors in our final specification to be conservative and to not bias our results in favor of credit specific (frailty) factor effects.

Allowing for one frailty factor is standard in the literature, see for example Duffie et al. (2009) and Azizpour et al. (2010), and appears sufficient to capture deviations of credit from macro conditions when modeling defaults. Since we have three different regions, we allow for three region-specific frailty factors. Finally, we allow for one industry-specific factor for financial firms. The latter loads on financial firms from all regions but may do so possibly to different degrees. The three macro-financial factors f_t^m are common to all macro covariates, thus taking into account the positive correlation between, say, U.S. and European conditions that may result from macroeconomic linkages and financial cross-holdings.

We refer to the Appendix A2 for further diagnostic checks for the chosen model specification. The diagnostic checks suggest that, while restrictive, this specification is sufficiently flexible to accommodate most of the heterogeneity observed across regions and industries for the purpose at hand.

Table 3: **Parameter estimates**

We report the maximum likelihood estimates of selected coefficients in the specification of the log-odds ratio (9) with parameterization (12) for $\lambda_{r,j}$ and $\beta_{r,j}$. Coefficients $\lambda_{r,j}$ combine to fixed effects, or baseline failure rates. Factor loadings $\beta_{r,j}$, $\gamma_{r,j}$, and δ_r refer to three macro factors f_t^m , three region-specific frailty factors f_t^d , and a financial industry specific risk factor f_t^i , respectively. The macro factors are common to all macro and default data across regions. As a result, U.S. macroeconomic conditions may affect the E.U. area and the Asia-Pacific (A.P.) region, and vice versa. The frailty factors load on financial and non-financial firms' defaults in a respective region. The estimation sample is from 1980Q1 to 2012Q2. The model has $k = 1, \dots, 3$ macro factors, three default risk specific factors (one for each region), and an industry-specific factor for the financial industry.

Intercepts			Loadings macro factor $f_{k,t}^m, k = 1, \dots, 3$			Loadings frailty and industry factor f_t^d, f_t^i		
$\lambda_{r,j} = \lambda_0^c + \lambda_{1,j}^c + \lambda_{2,r}^c + \lambda_{3,s}^c$			$\beta_{k,r,j} = \beta_{k,0}^c + \beta_{k,1,j}^c + \beta_{k,2,r}^c$			$\gamma_r = \gamma_{0,r}^c$		
par	val	t-val	par	val	t-val	par	val	t-val
λ_0^c	-4.66	10.12	$\beta_{1,0}^c$	-0.27	4.22	$\gamma_{0,US}^c$	0.35	5.32
			$\beta_{1,1,fin}^c$	0.29	4.73	$\gamma_{0,EU}^c$	0.45	2.52
$\lambda_{1,fin}^c$	0.00	0.02	$\beta_{1,2,EU}^c$	0.07	0.71	$\gamma_{0,AP}^c$	0.87	2.93
$\lambda_{2,EU}^c$	-0.32	1.31	$\beta_{1,2,AP}^c$	0.22	1.20			
$\lambda_{2,AP}^c$	-0.82	1.72						
$\lambda_{3,IG}^c$	-4.18	5.72	$\beta_{2,0}^c$	0.21	5.36	$\delta_r = \delta_{0,US}^c + \delta_{1,r}^c$		
			$\beta_{2,1,fin}^c$	-0.10	2.80	$\delta_{0,US}^c$	0.28	4.15
			$\beta_{2,2,EU}^c$	0.09	1.79	$\delta_{1,EU}^c$	0.17	1.41
			$\beta_{2,2,AP}^c$	0.20	1.95	$\delta_{1,AP}^c$	0.75	1.87
			$\beta_{3,0}^c$	-0.15	2.40			
			$\beta_{3,1,fin}^c$	-0.03	0.59			
			$\beta_{3,2,EU}^c$	0.16	1.83			
			$\beta_{3,2,AP}^c$	-0.04	0.25			

4.2 Financial sector systematic risk: parameter estimates and variance decomposition

Observed default fractions and EDF data for financial sector firms suggest that systematic default rates are up to ten times higher in bad times (say, 1988-90 and 2008-09 in the U.S.) than in good times (say, 1996-97 and 2005-06 in the U.S.). This is striking. Why do financial firm defaults cluster so dramatically over time? We contribute to this question by relating financial sector systematic risk to its underlying latent risk drivers, and investigating which factors contribute how much in terms of explanatory power.

The parameter estimates in Table 3 indicate that macro, frailty, and industry effects are

Table 4: Variance decomposition into macro, frailty, and industry effects

We report the results of a variance decomposition of transformed systematic default rates for financial sector firms in three economic regions. The unconditional variation in $\theta_{r,j,t} = \log(\pi_{r,j,t}) - \log(1 - \pi_{r,j,t})$ is attributed to three sources of stress. Each source of stress is captured by a corresponding set of latent factors and associated risk factor standard deviations. The respective variance shares are $s_{r,j}^m = \beta'_{r,j}\beta_{r,j}/\text{Var}(\theta_{r,j,t})$, $s_{r,j}^d = \gamma'_{r,j}\gamma_{r,j}/\text{Var}(\theta_{r,j,t})$, and $s_{r,j}^i = \delta'_{r,j}\delta_{r,j}/\text{Var}(\theta_{r,j,t})$, where $\text{Var}(\theta_{r,j,t}) = \beta'_{r,j}\beta_{r,j} + \gamma'_{r,j}\gamma_{r,j} + \delta'_{r,j}\delta_{r,j} \geq 0$, where r, j pick the appropriate cross section of financial firms. The estimation sample is from 1980Q1 to 2012Q2.

	Changes in observed macro-financial conditions	Latent default- specific dynamics	Latent financial sector dynamics
	$s_{r,fin}^m$	$s_{r,fin}^d$	$s_{r,fin}^i$
U.S.	18%	50%	32%
E.U. area	11%	44%	45%
Asia-Pacific region	10%	38%	52%

all important for international credit risk and financial sector risk conditions. Defaults from all regions and industries load on common factors from global macro-financial data (see the $\beta_{k,r,j}$ coefficients for macro factor $k = 1, \dots, 3$). This by itself already implies a considerable degree of default clustering. Statistical significance may not be high due to defaults being rare events. In general, however, common variation with macro data is not sufficient. Frailty effects are also found to be important in all regions ($\gamma_{r,j}$ coefficients). Moreover, the financial industry-specific factor also loads significantly on data from all regions (δ_r coefficients). The large amount of common variation in risk dynamics for financial firms across regions is intuitive. These groups operate globally both in terms of their lending and funding activities, and are exposed to highly interdependent financial markets around the globe.

Table 4 attributes the variation in the (Gaussian) log-odds of financial sector failure rates to three primary risk drivers, i.e., changes in macro-financial conditions, frailty risk for all firms (financial and non-financial), and financial sector-specific dynamics. These drivers are associated with the vectors of latent factors f_t^m , f_t^d , and f_t^i , respectively. The relative importance of each source of variation can be inferred from the estimated risk factor loadings. Given that each risk factor is unconditional standard normal, the factor loading is the estimated risk factor volatility (standard deviation) by construction.

Table 4 indicates that industry-specific variation and frailty dynamics are the most important sources of financial sector systematic default risk. This is likely due to a particularly important role for contagion and network effects in this sector. Changes to joint macro-financial and default conditions are not a dominant driver of financial distress; macro factors f_t^m are found to explain less than 20% of the total variation in financial sector systematic default risk. As an important caveat, the factor loading parameters are estimated with little precision from rare data on actual financial sector defaults (111 in the U.S., 20 in the E.U., and 16 in Asia-Pacific sub-samples). In addition, the variance shares are sensitive to changes in model specification and parameter pooling restrictions. We can conclude, however, that all three sources of risk, in particular frailty and industry-specific dynamics, should be accounted for in an overall risk assessment framework.

4.3 Tracking global financial sector stress over time

This section reports and discusses our in-sample empirical estimates of unobserved financial sector stress based on the the risk indicators defined in Section 2.3. We distinguish financial sector firms located in the U.S., Europe (E.U. area), and Asia-Pacific region as explained in Section 3.

Figure 3 plots model-implied financial sector default rates as in equation (13). The model-implied default rate is the share of overall intermediaries that can be expected to default over the next three months. Multiplying this rate by four gives an approximate annual rate. High probabilities of default are visible in the left panel for the U.S. during the recession years of 1991, 2001, and acute financial crisis from 2008-10. For European financial firms, the euro area debt crisis from 2010-12 at the end of the sample is particularly stressful. Model-implied stress for European intermediaries is lower than for U.S. financials for most of the sample pre-2010. This is partly due to sample composition, as rated European financials that access the capital markets tend to have a high credit rating on average, and observed historical default frequencies are low.

Figure 3 further compares the model-implied default rate (solid line) with a mean ex-

pected default frequency (EDF) estimate from Moody's KMV (dashed line) and ex-post realized quarterly default fractions (symbols). The Moody's database contains both rated banks as well as non-bank financial firms, such as insurers and real estate firms. Our Moody's KMV rate estimate refers to about twenty large financial firms and thus has a much more limited base.

The simultaneous default of multiple financial sector firms is a relatively rare event, but with potentially a large impact if the event materializes. The top three panels of Figure 4 plot the probability of at least $k\%$ financial firms failing over a one year horizon (vertical axis), as a decreasing function of k . The model implied probabilities are computed over the period 1985Q1 to 2012Q2 (horizontal axis). The three-dimensional graphs roughly follow what is implied by the aggregated failure rate for the financial sector as plotted in Figure 3. The three bottom panels of Figure 4 cut the three-dimensional plots into slices along the time dimension, at 0.5%, 1%, and 2% of overall financial sector firms. For example, in the E.U. area at the end of the sample, the probability of failure of at least 1% of financial sector firms (e.g., at least four firms of average size out of roughly four hundred firms in the Moody's database), at coincident levels of stress, is around 50%. This is a substantial risk.

Finally, Figure 5 plots financial distress based on indicator (15). The probit transform in (15) maps model-implied systematic stress into a uniform random variable, such that its percentiles can be read off from the transformed y-axis. Financial distress is highest towards the end of the sample (above the 80th percentile), and virtually absent during the late-1990s and mid-2000s (below the 20th percentile). The dotcom bubble in the late 1990s and the housing bubble from approximately 2005-2007 have started to build up during such periods of low risk. This is in line with systemic risk buildup that occurs during (exceptionally) good times when *measured* risks are low, as suggested by Borio (2010, 2011), Boissay, Collard, and Smets (2013), and Brunnermeier and Sannikov (2013).

Figure 3: Implied default rate for financial sector firms

The panel plots the estimated default hazard rate $\pi_{r,j,t}$ for financial sector firms in the United States, European Union area, and Asia-Pacific region, respectively. The panels compare the model-implied rates with the mean expected default frequencies for a smaller number of large financial firms in each region, and observed default fractions from the default count data panel. The estimate is a quarterly rate, multiplying by four gives an approximate annual rate. The estimation sample is from 1980Q1 to 2012Q2.

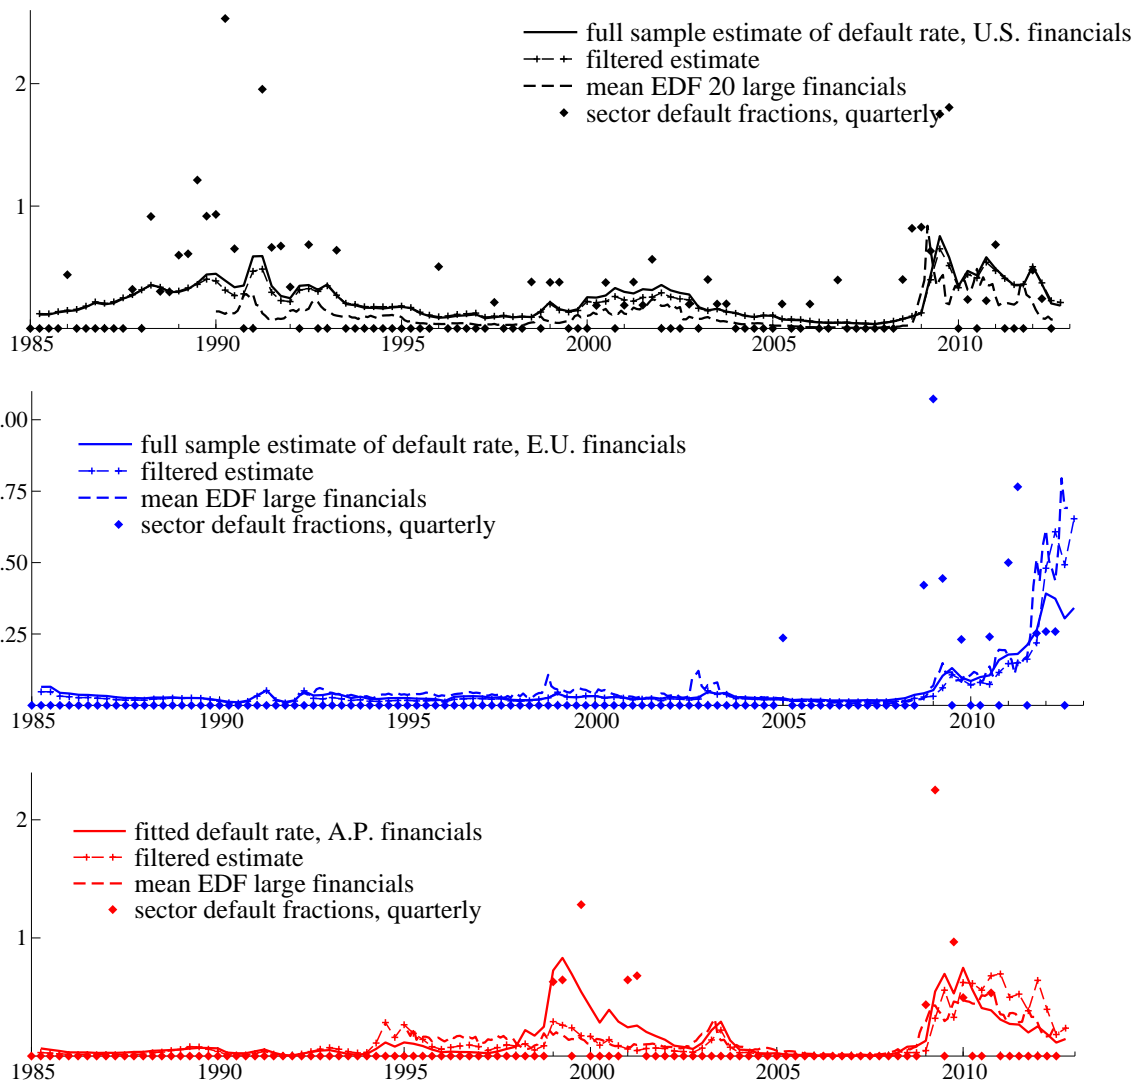


Figure 4: Probability of financial sector firm joint defaults

The top three graphs plot the probability of a simultaneous failure of $k\%$ or more financial firms over a one year period in the U.S., the E.U. area, and the Asia-Pacific region, respectively. The horizontal axis measures time from 1985Q1 to 2012Q2. The vertical axis measures the time-varying probability as a decreasing function of k . The bottom three panels report slices through the above three-dimensional plots along the time dimension, and refer to financial firms headquartered in the U.S., the E.U., and the Asia-Pacific region, respectively.

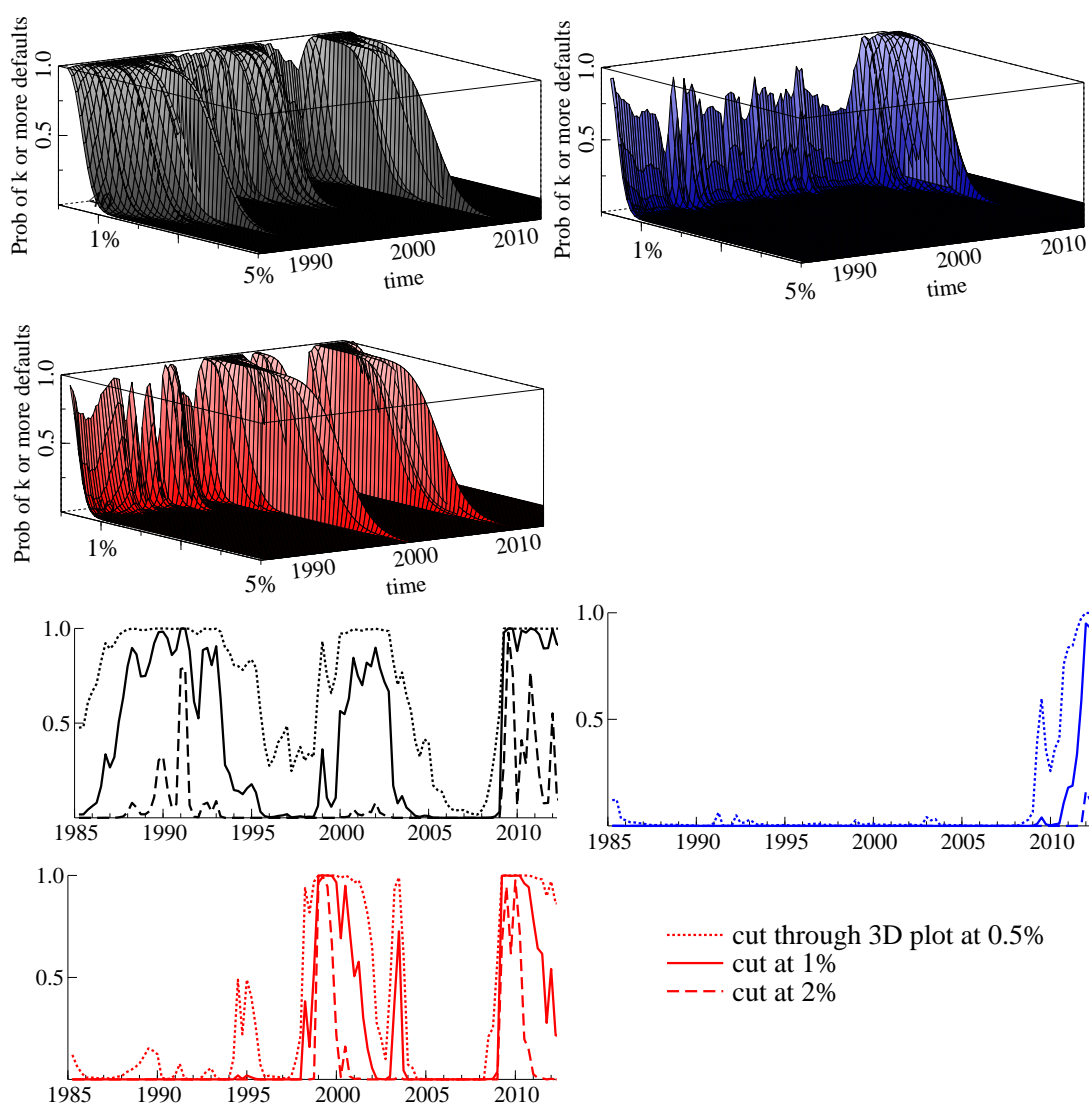
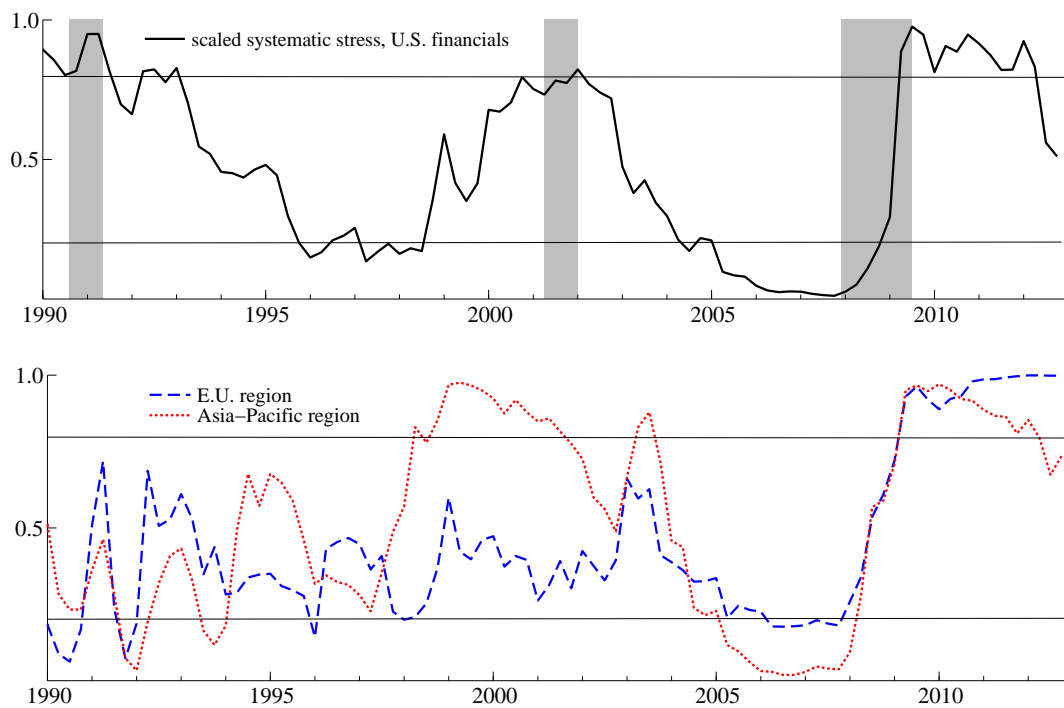


Figure 5: PIT transform of financial sector systematic default rate

The figure plots the probability integral transform (15) for financial firm systematic default risk in the U.S. (top panel) and the E.U. and Asia-Pacific region (bottom panel). The probit transformation of the log-odds ratio implies that percentiles of systematic stress can be read off the y-axis. Two horizontal lines at 20% and 80% indicate when estimated systematic stress for financial firms is particularly low and particularly high, respectively. Light shaded areas in the top panel indicate U.S. recession periods according to NBER.



4.4 Nowcasts and forecasts of global financial sector stress

In this section we assess the now-casting and out-of-sample forecasting performance of our favorite model specification. We compare model-based now- and forecasts to competing now- and forecasts from two main modeling alternatives: an exponentially weighted moving average specification of the default rate, as well as the mean expected default frequency (EDF) for financial firms as obtained from Moody's KMV. Accurate forecasts of financial sector stress are valuable for counterparty credit risk management, conditional stress testing exercises, as well as financial sector surveillance from a prudential perspective. 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 based on data from 1980Q1 to 2007Q4. Now-casts and forecasts of quarterly default rates from 2008Q1 to 2012Q2 are computed conditional on these training sample parameter values. The model-based predictions are

$$\begin{aligned}\hat{\pi}_{r,j,t}^{\text{Model, now}} &= \left(1 + e^{-\hat{\theta}_{r,j,t}^{\text{Model, now}}}\right)^{-1}, \\ \hat{\pi}_{r,j,t+1}^{\text{Model, forc}} &= \left(1 + e^{-\hat{\theta}_{r,j,t+1}^{\text{Model, forc}}}\right)^{-1},\end{aligned}$$

where $\hat{\theta}_{r,j,t}^{\text{Model, now}}$ and $\hat{\theta}_{r,j,t+1}^{\text{Model, forc}}$ are based on (9), where risk factors are either filtered factor estimates $f_{t|t}$ or predicted one-step ahead estimates $f_{t+1|t}$, respectively.

We consider two main alternatives for the now-casting of current-quarter default rates and forecasting them 1-quarter ahead. First, we fit a standard exponentially weighted moving average (EWMA) to quarterly observed default fractions $\hat{\pi}_{r,j,t} = y_{r,j,t}/k_{r,j,t}$.

$$\begin{aligned}\hat{\pi}_{r,j,t}^{\text{EWMA, now}} &= \alpha \hat{\pi}_{r,j,t} + (1 - \alpha) \hat{\pi}_{r,j,t-1}^{\text{EWMA, now}}, \\ \hat{\pi}_{r,j,t+1}^{\text{EWMA, forc}} &= \alpha \hat{\pi}_{r,j,t} + (1 - \alpha) \hat{\pi}_{r,j,t}^{\text{EWMA, forc}},\end{aligned}$$

where the smoothing parameter α in the EWMA specification is chosen to give the closest mean squared error fit to the target rate in the training sample. As a second forecasting alternative, we consider the current mean expected default frequency (EDF) for large financial sector firms as a forecast. To this purpose the current annual EDF is converted into a

Table 5: Nowcasts and out-of-sample forecasts of regional financial sector stress

We report mean absolute forecast error (MAE) and root mean squared forecast error (RMSE) statistics associated with now-casts and out-of-sample forecasts of financial sector systematic default rates for three economic regions. The estimation sample is from 1980Q1 to 2007Q4; now-casts and out of sample forecasts are obtained for the financial crisis period for 18 quarters from 2008Q1 to 2012Q2. Now-casts are obtained from filtered risk factors, out of sample forecasts are obtained from one-quarter ahead predicted factors.

Error statistic		MAE	RMSE	MAE	RMSE
Target probability / rate:		full model smoothed		EDF large financials	
US fin	Model nowcast	0.024%	0.026%	0.610%	0.786%
	Model nowcast as forecast	0.108%	0.133%	0.623%	0.844%
	Model forecast 1q ahead	0.023%	0.026%	0.610%	0.786%
	EWMA nowcast	0.286%	0.323%	0.912%	1.126%
	EWMA forecast	0.327%	0.363%	0.953%	1.165%
	EDF-based forecast	0.632%	0.811%		
EU fin	Model nowcast	0.004%	0.007%	0.511%	0.711%
	Model nowcast as forecast	0.043%	0.077%	0.538%	0.761%
	Model forecast 1q ahead	0.005%	0.008%	0.510%	0.710%
	EWMA nowcast	0.149%	0.229%	0.664%	0.944%
	EWMA forecast	0.143%	0.224%	0.656%	0.938%
	EDF-based forecast	0.515%	0.718%		
AP fin	Model nowcast	0.288%	0.420%	0.589%	0.634%
	Model nowcast as forecast	0.380%	0.493%	0.613%	0.761%
	Model forecast 1q ahead	0.288%	0.420%	0.590%	0.634%
	EWMA nowcast	0.336%	0.398%	1.152%	1.308%
	EWMA forecast	0.369%	0.432%	1.190%	1.345%
	EDF-based forecast	0.822%	0.916%		

quarterly rate as

$$\hat{\pi}_{r,j,t+1}^{\text{EDF, forc}} = 1 - (1 - \text{EDF}_{r,j,t})^{\frac{1}{4}},$$

where $\text{EDF}_{r,j,t}$ is a one-year ahead default risk estimate.

The measurement of forecasting accuracy of time-varying default probabilities is not entirely straightforward. Observed default fractions are only a crude measure of default conditions. We can illustrate this inaccuracy by considering a group of, say, 10 firms. Even if the default probability is forecast perfectly, it is unlikely to coincide with the observed default fraction of either 0, 1/10, 2/10, etc. The forecast error may therefore be large but it does not necessarily indicate a bad forecast. We therefore consider two benchmarks that do not suffer from this drawback. First, we use the model-implied financial sector default rate, based on the full-sample parameter and risk factor estimates. Second, we consider the mean EDF rate for large financial institutions as an alternative benchmark rate. Both benchmark rates are graphed in Figure 3.

Table 5 summarizes the results from our forecasting study. We make two main observations. First, the model-implied now-casts (based on filtered risk factors) and one-quarter ahead out-of-sample forecasts (based on out-of-sample predicted risk factors) are consistently the entries with the lowest mean absolute error (MAE) and root mean squared error (RMSE) statistics. Our modeling framework produces more accurate forecasts than the alternative EWMA-based specification and the EDF-based forecast. This is the case for the U.S., and also in the European and Asia-Pacific data subsamples. The model now- and forecasts continue to be the most accurate even if forecasts are compared with EDF rates as benchmark rates instead of model-implied full sample estimated rates. Second, the out-of-sample forecasts of one-quarter ahead default rates are more accurate than when the now-casts for the current quarter are used to predict default rates one step ahead. This suggests that the estimated risk factor dynamics capture salient and robust features of co-movement in financial sector default data over time. Together, these observations suggest that effectively combining different data inputs - EDFs, default fractions, and macro-financial data - improves inference on the unobserved quantity of interest, i.e., current quarter and projected future

values of financial sector systematic default risk.

4.5 Credit risk dislocations as a warning signal

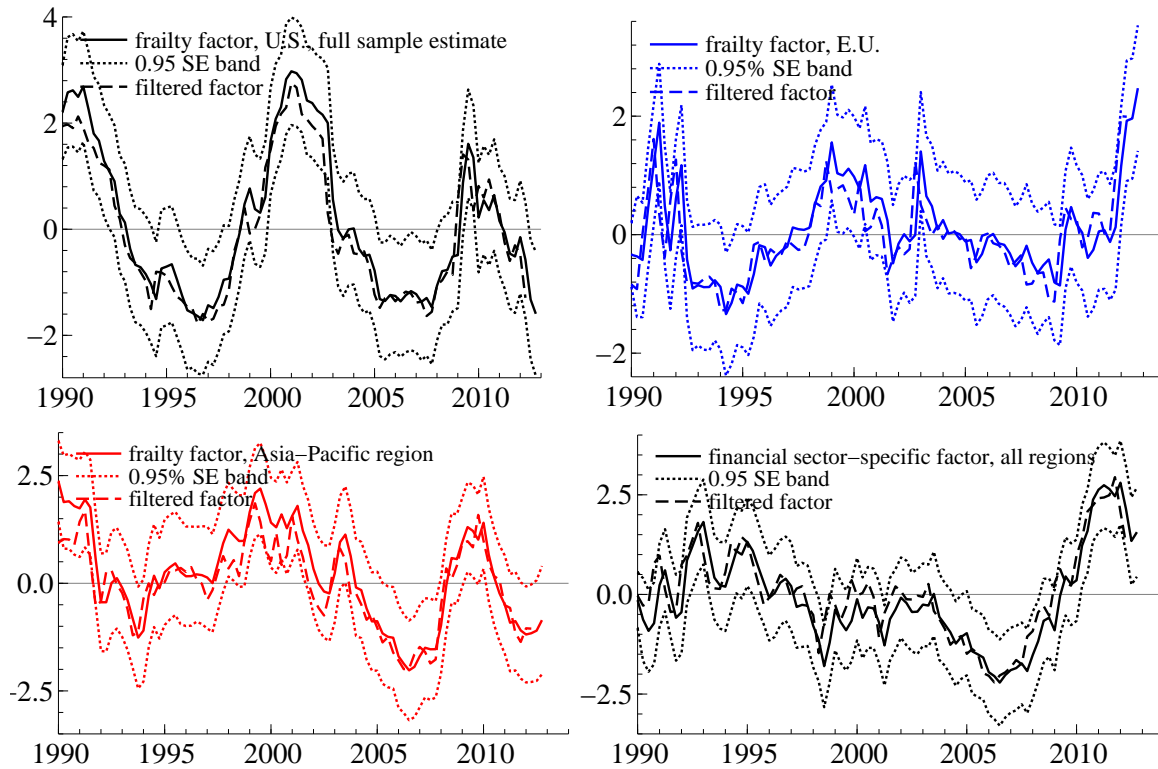
The coincident indicators presented in the previous section are helpful ingredients in a prudential risk monitoring toolbox. However, coincident indicators are not constructed to anticipate crises. The unobserved factor estimates obtained from our model, however, can be used to construct alternative measures that can be used for early warning purposes.

Our warning signal builds on Duffie, Eckner, Horel, and Saita (2009), Azizpour, Giesecke, and Schwenkler (2010), Koopman, Lucas and Schwaab (2011, 2012), and Creal et al. (2013) who all find substantial evidence for a dynamic unobserved risk factor driving U.S. firm defaults above and beyond what is implied by observed macro-financial covariates and other information. We interpret the credit risk specific or frailty factor as largely capturing unobserved variation in credit supply, or changes in the ease of credit access. We rely on two pieces of evidence for interpretation, as reported in Koopman, Lucas, and Schwaab (2011, 2012). First, frailty tends to load more heavily on financially weaker – and thus more credit constrained – firms. This appears to hold in general, and in particular during the years leading up to the financial crisis. Second, our frailty estimates are highly correlated with ex post reported lending standards, such as the ones obtained from the Senior Loan Officer Survey (SLO) and as reported in Maddaloni and Peydro (2011). Clearly, credit risks and credit quantities are related: it is hard to default if firms and households are drowning in credit. Conversely, even solvent firms may come under stress if credit is rationed at the economy wide level. As a result, systematic default risk conditions (the default cycle) can decouple from what is implied by the macro-financial environment (the business cycle).

Tracking credit *risk* conditions and their deviations from macro-financial fundamentals is related to the notion of tracking credit *quantities* (or aggregates) over time. The usefulness of the private-credit to GDP-ratio as an early warning indicator for costly asset price boom and busts and systemic banking crises is a recurring finding in the early warning literature, see for example Borio and Drehmann (2009) and Alessi and Detken (2011). We argue that

Figure 6: Latent risk factor estimates

The left and middle panels report the conditional mean estimates of the frailty factor for U.S. and E.U. corporates, respectively. The right panel plots the financial sector specific risk factor. The latter loads on financial firms in all regions. The approximate standard error bands are plotted at a 95% confidence level and refer to the full sample estimates. The reported factor estimates are from 1990Q1-2012Q2 because data up to 1990 is sparse. Filtered factor estimates are plotted for comparison.



credit quantities and risk are related - virtually no firm defaults during a serious lending bubble - and that an analysis of credit risk conditions over time can therefore give a valuable complementary perspective on credit market activity.

The left panel of Figure 6 presents the estimated frailty factors for the U.S., European Union area, and Asia-Pacific region. 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 lower than would be expected from observed macro and financial data.

At these times, frailty effects are negative, and can be significantly different from zero. On top of these developments, our estimate of the financial industry-specific risk factor indicates a particularly benign risk environment for financial firms during these years leading up to the crisis.

The top and bottom panels of Figure 7 plot ‘credit risk deviations’ (CRD) as specified in (16) for financial and non-financial firms, respectively. The indicator combines estimated frailty and financial industry effects into a warning signal. The indicator captures the extent to which local stress in a given industry sector differs from what macro-financial fundamentals would suggest. The figure compares estimated deviations in the U.S., the E.U. area, and Asia-Pacific region. Light and dark shaded areas correspond to NBER recession periods for the U.S. and times of banking crises as identified in Laeven and Valencia (2010), respectively. The graph is based on filtered risk factor estimates that take into account information which is available at time t .

We indicate three banking crises in the graphs, two in the U.S. (1988 and 2007-10) and one in the E.U. area (2008-10). Figure 7 draws signal thresholds at a 90% confidence level. The figure reveals a particularly large and persistent decoupling of risk conditions from fundamentals for both financial and non-financial firms preceding the financial crisis and recession of 2007-2009. Here, risk conditions were significantly and persistently below what was suggested by fundamentals. Such a development may indicate a lending bubble, in particular if credit quantity growth is unusually high as well and bank lending standards are generous (which has been the case). Conversely, the indicator may also signal risk conditions that are considerably worse than suggested by fundamentals. Examples are the U.S. conditions during the years 1988-90 leading up to the 1991 recession, and conditions during 2010 to 2012Q2 in all three regions. For financial firms, the end of the sample marks the start of the Euro area sovereign debt crisis. These developments may not be captured by the macro fundamentals and are therefore absorbed by the latent industry factor, see Figure 6. As such, they become part of the credit risk deviations indicator in Figure 7.

We now investigate the usefulness of our credit risk deviations indicator for forward

Figure 7: Scaled credit risk deviations (CRD)

We plot filtered deviations of credit risk conditions from macro-financial fundamentals as measured in (16). “Filtered” means that only information available at time t is used for signal extraction. The top and bottom charts refer to financial firms’ and ‘non-financial firms’ risk conditions, respectively. Light and dark shaded areas correspond, respectively, to NBER U.S. recession periods and times of systemic banking crises as identified in Laeven and Valencia (2010). The indicator is a standard normal variable by construction; the horizontal lines indicate an approximate 90% confidence interval.

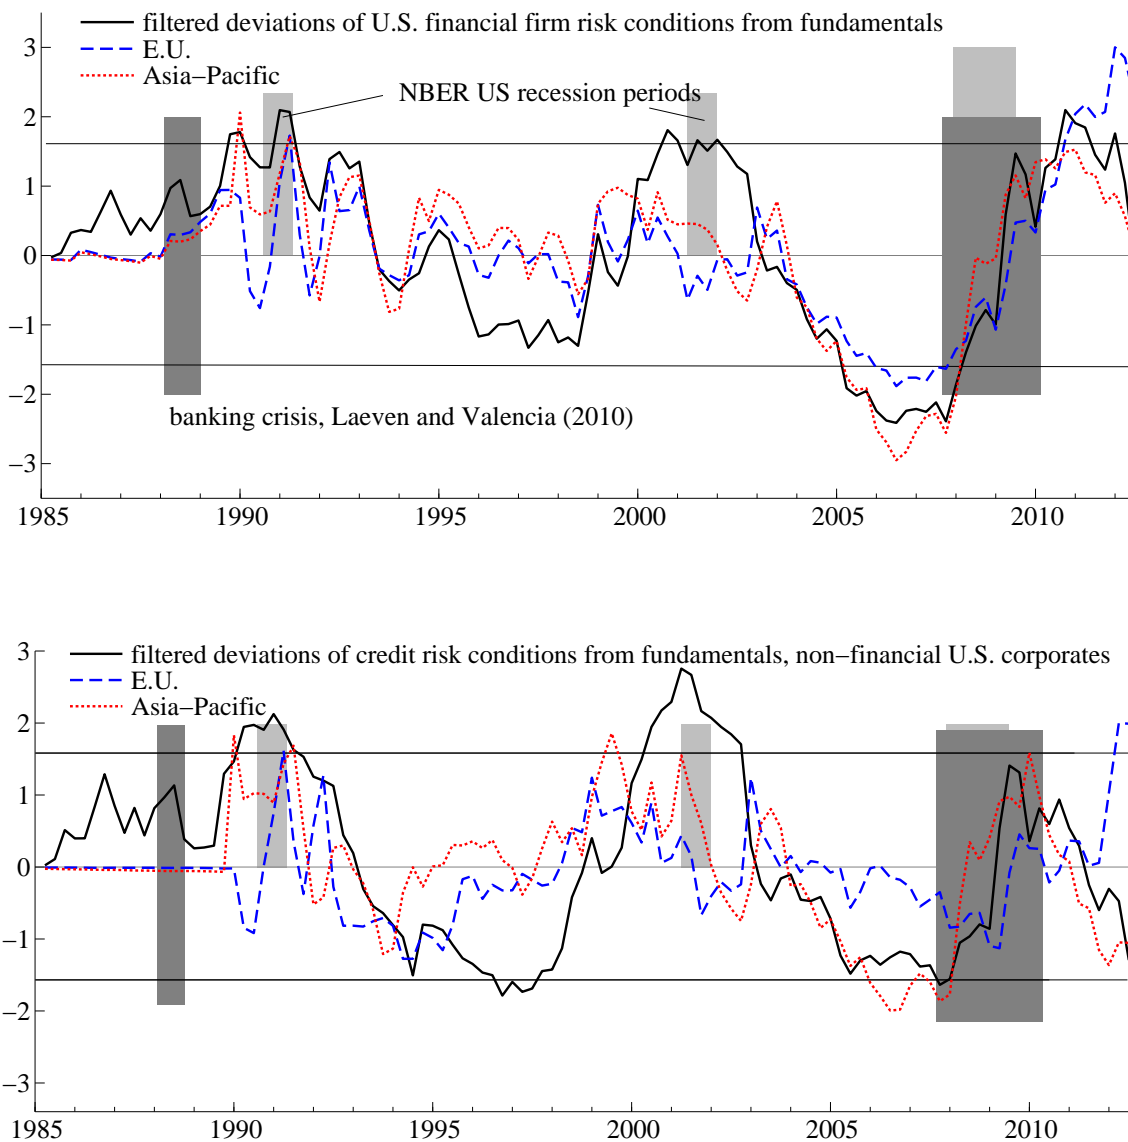


Table 6: Ranking of credit based early warning indicators

We rank several credit based early warning indicators in terms of their usefulness in predicting credit fueled asset price booms that later bust and turn out to be costly in real terms. The data, definitions, and methodology are identical to Alessi and Detken (2011). Economic loss is defined as a preference weighted average of type-1 and type-2 errors, $\text{loss} = \theta C / (A + C) + (1 - \theta) B / (B + D)$, where letters A to D indicate the number of quarters and countries for which A: a signal is issued and coincides with a costly asset price boom, B: a signal is issued but does not coincide with a costly boom, C: no signal coincides with a costly boom, D: no signal is issued and coincides with no costly boom, and θ is a preference parameter, which is here set to 0.4, implying a slightly higher aversion against false warnings. The optimal signal threshold is chosen as a multiple of 0.05 and corresponds to the percentile of the indicator's distribution that achieves minimal loss. The 'signal' is $A / (A + C)$ and indicates the share of correctly classified costly booms. The share of type I errors is $1 - A / (A + C) = C / (A + C)$, i.e., no signal is issued but a costly boom occurs, as a share of quarters for which a costly boom occurs. The fraction $B / (B + D)$ is the share of type II errors, i.e., a warning is issued but no boom occurs, as a share of quarters that are non-booms. The probability of being in a costly boom given a signal is given by $A / (A + B)$. The noise-to-signal ration (N2S) is standard in the early warning literature and is computed as $[B / (B + D)] / [A / (A + C)]$. The CRD indicators are mirrored at the horizontal axis, such that large negative deviations trigger a signal. The test sample consists of 11 European countries from 1985Q1 to 2008Q4. Data is only available for 1998Q1 to 2008Q4 for the indicators in the last two rows.

Indicators	loss	opt signal threshold	booms called $A / (A + C)$	type-2 err $B / (B + D)$	Pr[boom signal] $A / (A + B)$	N2S ratio
Global PC gap	0.35	65	57%	30%	32%	0.53
CRD fin EU	0.36	90	23%	8%	41%	0.34
CRD fin US	0.36	90	21%	7%	43%	0.32
PC/GDP gap	0.36	70	48%	25%	32%	0.51
Loan/Deposits gap	0.42	90	6%	8%	21%	1.30
Total Assets/GDP gap	0.42	95	3%	4%	18%	1.50

looking early warning purposes. Table 6 ranks several credit based early warning indicators in terms of their usefulness in predicting credit fueled asset price booms that later bust and are costly in real terms. The data, definitions, and methodology are identical to Alessi and Detken (2011). The asset price booms refer to a weighted average of residential property and equity prices and are based on BIS data. By focusing on such asset prices, it is much easier to define a boom and bust for a cross section of countries than when focusing on banking stress directly. The test sample consists of 11 European countries (Belgium, Germany, Denmark, Spain, Finland, France, the U.K., Ireland, the Netherlands, Norway, and Sweden), from 1985Q1 to 2008Q4.

According to our analysis, the global private credit to GDP ratio gap remains the best indicator for early warning purposes. Economic loss, defined as a preference-weighted average of type I (missing a crisis) and type II errors (false warning), is lowest for this indicator. The optimal signal threshold, however, is relatively low, which means that many signals are issued. Many costly booms are called, but the number of type II errors is also high.

Our credit risk deviations indicators (16) for E.U. and U.S. financial firms come in as a close second. The optimal signal threshold is higher (90th percentile), which means that less signals are issued than for the global PC gap. Less booms are called (23% instead of 57%), but the type II error is also much lower (8% instead of 30%). The noise-to-signal ratios are the two lowest entries and suggest that the indicator is fairly reliable in this sense. Our credit risk based indicator does better than many alternative credit quantity-based measures, such as the private credit to GDP ratio gap based on individual countries' time series, the loan-to-deposits ratio, and total assets to GDP.

Finally, the pronounced deviations of risk conditions from fundamentals are relatively robust to variations in the set of explanatory right-hand side variables. For example, Duffie, Eckner, Horel, and Saita (2009) report that pre-crisis U.S. frailty effects cannot easily be attributed to omitted standard covariates such as real GDP growth or industrial production, which have been missing in their analysis. Koopman, Lucas, and Schwaab (2011) include more than 100 macro-financial covariates in their empirical study of U.S. data, and still

find important frailty effects. As a result, default and business cycle activity appear to be imperfectly related. The extent to which they have decoupled can be captured by our modeling framework and used to flag exceptional credit market developments.

5 Conclusion

We formulated and estimated a nonlinear, partly non-Gaussian panel time series model for international credit risk and macro linkages. Based on this new model, we proposed a novel framework for aggregate financial sector risk assessment. We combined the risk factor and parameter estimates into new and straightforward coincident and forward looking indicators of systematic financial sector risk. Methodologically, the factor structure in our modeling framework allows us to address the computational challenges associated with a large cross sectional dimension of financial firms across the globe. In addition, our framework easily allows us to combine different sets of panel data in a single integrated framework. We find that the measurement of credit *risk* conditions over time yields an important complementary perspective on credit market activity to the measurement of credit *quantities*. A decoupling of credit risk conditions from macro-financial fundamentals may indicate too loose lending standards and too easy credit access. In this way, it may serve as an early warning signal for prudential policy makers. In an extensive comparison, such new credit risk related measures compare favorably to other suggested risk indicators, such as the private credit to GDP gap.

References

- Acharya, V. V., L. H. Pedersen, T. Philippon, and M. Richardson (2010). Measuring systemic risk. *NYU working paper*.
- Adrian, T. and M. Brunnermeier (2009). Covar. *Federal Reserve Bank of New York Staff Reports* 348.
- Alessi, L. and C. Detken (2011). Quasi real time early warning indicators for costly asset price boom/bust cycles: a role for global liquidity. *European Journal of Political Economy* 27(3),

520–533.

Aoki, K. and K. Nikolov (2012). Bubbles, banks and financial stability. *ECB working paper 1495*, 1–52.

Azizpour, S., K. Giesecke, and G. Schwenkler (2010). Exploring the sources of default clustering. *Stanford University working paper*.

Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica* 70(1), 191–221.

Bharath, S. T. and T. Shumway (2008). Forecasting Default with the Merton Distance to Default Model. *The Review of Financial Studies* 21(3), 1339–1369.

Billio, M., M. Getmansky, A. W. Lo, and L. Pelizzon (2012). Econometric measures of systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, forthcoming.

Boissay, F., F. Collard, and F. Smets (2013). Booms and Systemic Banking Crises. *ECB working paper 1514*, 1–55.

Borio, C. (2010). Implementing a macroprudential framework: Blending boldness and realism, Keynote address for the BIS-HKMA research conference on Financial Stability: Towards a Macroprudential Approach. Available at www.bis.org. *Hong Kong SAR, 5-6 July 2010*.

Borio, C. (2011). Rediscovering the macroeconomic roots of financial stability policy: journey, challenges and a way forward. *BIS Working Papers* 354, 1–34.

Borio, C. and M. Drehmann (2009). Assessing the risk of banking crises - revisited. *BIS Quarterly Review*, March 2009.

Brownlees, C. T. and R. Engle (2010). Volatility, correlation, and tails for systemic risk measurement. *NYU working paper*.

Brunnermeier, M. K. and Y. Sannikov (2013). A macroeconomic model with a financial sector. *American Economic Review*, forthcoming.

CEBS (2010). Committee of European Banking Supervisors: Summary Report from 23 July 2010. available at www.eba.europa.eu.

Creal, D., B. Schwaab, S. J. Koopman, and A. Lucas (2013). An observation driven mixed

- measurement dynamic factor model with application to credit risk. *Review of Economics and Statistics*, forthcoming 95.
- Das, S., D. Duffie, N. Kapadia, and L. Saita (2007). Common Failings: How Corporate Defaults Are Correlated. *The Journal of Finance* 62(1), 93–117(25).
- Duffie, D., A. Eckner, G. Horel, and L. Saita (2009). Frailty Correlated Default. *Journal of Finance* 64(5), 2089–2123.
- Duffie, D., L. Saita, and K. Wang (2007). Multi-Period Corporate Default Prediction with Stochastic Covariates. *Journal of Financial Economics* 83(3), 635–665.
- Durbin, J. and S. J. Koopman (1997). Monte Carlo Maximum Likelihood estimation for non-Gaussian State Space Models. *Biometrika* 84(3), 669–684.
- Durbin, J. and S. J. Koopman (2001). *Time Series Analysis by State Space Methods*. Oxford: Oxford University Press.
- Geweke, J. (1989). Bayesian inference in econometric models using Monte Carlo integration. *Econometrica* 57, 1317–39.
- Giesecke, K. and B. Kim (2011). Systemic risk: What defaults are telling us. *Management Science* 57, 1387–1405.
- He, Z. and W. Xiong (2012). Rollover risk and credit risk. *Journal of Finance* 67, 391–429.
- Huang, X., H. Zhou, and H. Zhu (2009). A framework for assessing the systemic risk of major financial institutions. *Journal of Banking and Finance* 33, 2036–2049.
- IMF (2009). Global financial stability report.
<http://www.imf.org/External/Pubs/FT/GFSR/2009/01/index.htm>.
- Koopman, S., R. Kräussl, A. Lucas, and A. Monteiro (2009). Credit cycles and macro fundamentals. *Journal of Empirical Finance* 16, 42–54.
- Koopman, S. J. and A. Lucas (2008). A Non-Gaussian Panel Time Series Model for Estimating and Decomposing Default Risk. *Journal of Business and Economic Statistics* 26 (4), 510–525.
- Koopman, S. J., A. Lucas, and A. Monteiro (2008). The Multi-Stage Latent Factor Intensity Model for Credit Rating Transitions. *Journal of Econometrics* 142(1), 399–424.

- Koopman, S. J., A. Lucas, and B. Schwaab (2011). Modeling frailty correlated defaults using many macroeconomic covariates. *Journal of Econometrics* 162(2), 312–325.
- Koopman, S. J., A. Lucas, and B. Schwaab (2012). Dynamic factor models with macro, frailty, and industry effects for u.s. default counts: the credit crisis of 2008. *Journal of Business and Economic Statistics* 30(4), 521–532.
- Laeven, L. and F. Valencia (2010). Resolution of Banking Crises: The Good, the Bad, and the Ugly. *IMF working paper 10/146*.
- Lando, D. and M. S. Nielsen (2009). Correlation in corporate defaults: Contagion or conditional independence. *Copenhagen Business School, working paper*.
- Maddaloni, A. and J.-L. Peydro (2011). Bank risk-taking, securitization, supervision and low interest rates - evidence from the Euro area and the U.S. lending standards. *Review of Financial Studies* 24, 2121–65.
- 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.
- Moreno, M. R. and J. I. Pena (2010). Systemic risk measures: the simpler the better? *UPF working paper*.
- 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.
- Segoviano, M. A. and C. Goodhart (2009). Banking stability measures. *IMF Working Paper WP/09/04*.
- Shephard, N. and M. K. Pitt (1997). Likelihood analysis of non-Gaussian measurement time series. *Biometrika* 84, 653–67.
- Stock, J. and M. Watson (2002). Forecasting Using Principal Components From a Large Number of Predictors. *Journal of the American Statistical Association* 97(460), 1167–1179.
- Tarullo, D. K. (2010). Equipping Financial Regulators with the Tools Necessary to Monitor Systemic Risk. *Statement before the Subcommittee on Security and International Trade and*

Appendix A1: estimation via importance sampling

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

$$p(y; \psi) = \int p(y, f; \psi) df = \int p(y|f; \psi)p(f; \psi)df, \quad (\text{A.1})$$

where $p(y|f; \psi)$ is the density of y conditional on f and $p(f; \psi)$ is the density of f . Importance sampling refers to the Monte Carlo estimation of $p(y; \psi)$ by sampling f from a Gaussian importance density $g(f|y; \psi)$. We can express the observation density function $p(y; \psi)$ by

$$p(y; \psi) = \int \frac{p(y, f; \psi)}{g(f|y; \psi)} g(f|y; \psi) df = g(y; \psi) \int \frac{p(y|f; \psi)}{g(y|f; \psi)} g(f|y; \psi) df. \quad (\text{A.2})$$

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

$$\tilde{p}(y; \psi) = g(y; \psi) M^{-1} \sum_{k=1}^M \frac{p(y|f^{(k)}; \psi)}{g(y|f^{(k)}; \psi)}, \quad f^{(k)} \sim g(f|y; \psi), \quad (\text{A.3})$$

is numerically efficient, see 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

$$y_{j,t} = \mu_{j,t} + \theta_{j,t} + \varepsilon_{j,t}, \quad \varepsilon_{j,t} \sim \text{N}(0, \sigma_{j,t}^2), \quad (\text{A.4})$$

where mean correction $\mu_{j,t}$ and variance $\sigma_{j,t}^2$ 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 $\mu_{j,t}$ and $\sigma_{j,t}$ 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 $g(f|y; \psi)$, a simulation smoother can be applied to the approximating model (A.4), see Durbin and Koopman (Ch. 11, 2001). For a set of M draws of $g(f|y; \psi)$, the evaluation of (A.3) 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 (5) and (4), density $g(y|f; \psi)$ is based on the Gaussian density for $y_{j,t} - \mu_{j,t} - \theta_{j,t} \sim N(0, \sigma_{j,t}^2)$, that is (A.4), and $g(y; \psi)$ can be computed by the Kalman filter applied to (A.4), see Durbin and Koopman (2001).

The likelihood function can be evaluated for any value of ψ . By keeping the random numbers fixed, we maximize the likelihood estimator (A.3) with respect to ψ by a numerical optimization method. 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 $\tilde{f}_t = E(f|y; \psi)$ and its standard error s_t via importance sampling can be achieved by

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

with $w_k = p(y|f^{(k)}; \psi)/g(y|f^{(k)}; \psi)$, $f^{(k)} \sim g(f|y; \psi)$, and \tilde{f}_t is the t th element of \tilde{f} .

Appendix A2: residual diagnostics for default panel

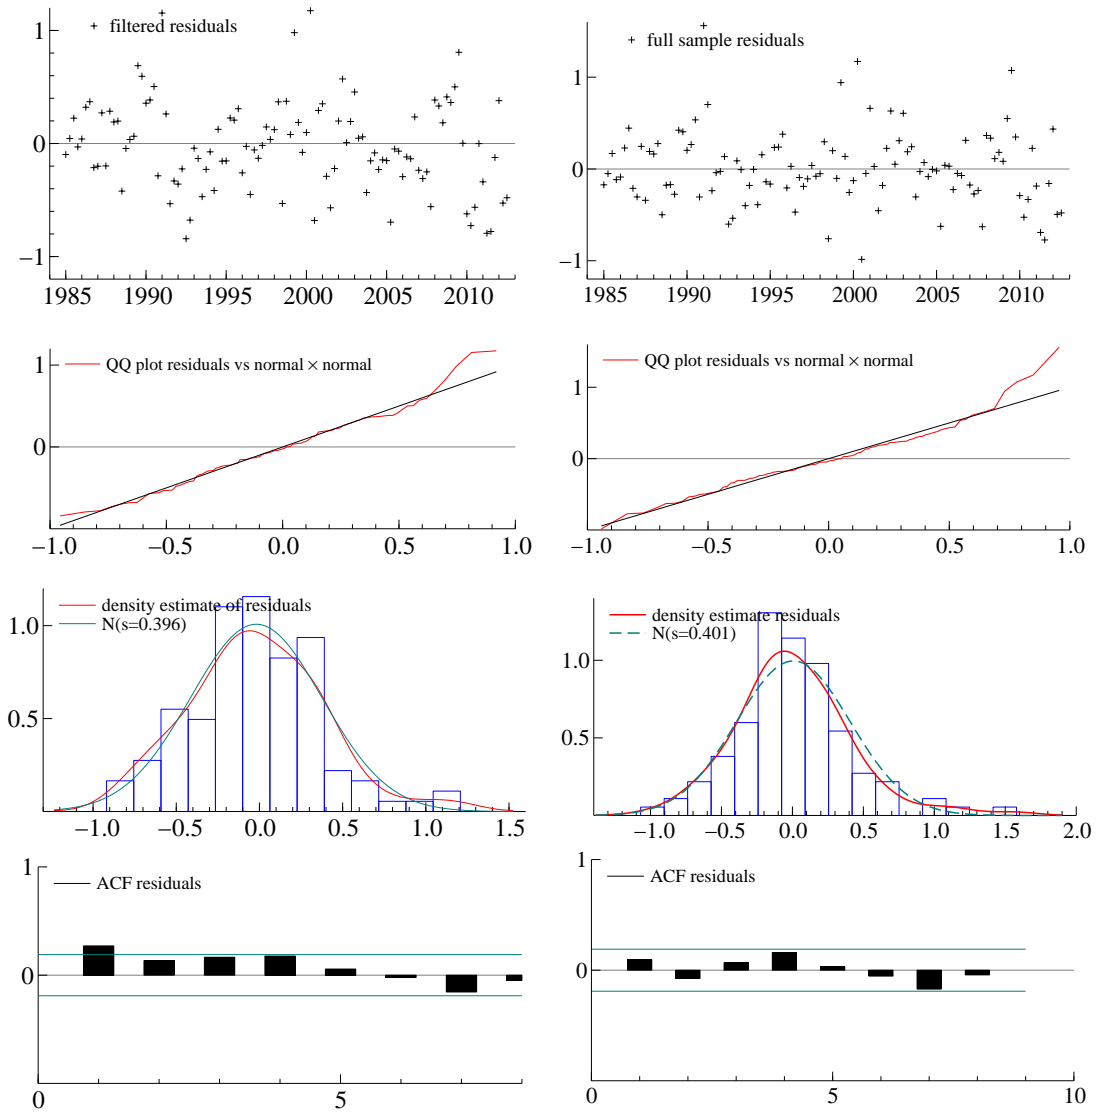
This section reports residual diagnostic checks for the non-Gaussian panel of international default counts. Pooling over parameters as in (12) and specified in Section 4.1 is restrictive, but keeps the model tractable and helps in estimating all model parameters and latent factors from sparse default data. Based on the residual diagnostics from this section, we conclude that our current parameter specification effectively combines model parsimony with the ability to test relevant hypotheses given the data at hand.

Similar to Koopman, Lucas, and Schwaab (2011), we define a time series of prediction errors as

$$\hat{e}_t = \left(\sum_{r,j} \delta_{r,j,t} \right)^{-1} \sum_{r,j} \delta_{r,j,t} (y_{r,j,t} - \hat{\pi}_{r,j,t} k_{r,j,t}), \quad (\text{A.5})$$

Figure 8: Diagnostic check for filtered and full sample estimates

We report error statistics that are based on filtered (on left hand side) and full sample estimates (right hand side) of default rates $\hat{\pi}_{r,j,t}$. We report, from top to bottom, (a) errors \hat{e}_t over time, (b) a QQ plot of prediction errors against the normal, (c) an error histogram and density estimate, (d) the error autocorrelation function.



where $y_{r,j,t}$ are observed quarterly default counts for firms of cross section j in region r at time t , $k_{r,j,t}$ are the respective number of firms at risk at the beginning of quarter t , the indicator function $\delta_{r,j,t}$ is equal to one if $k_{r,j,t} > 0$ and zero otherwise, and $\hat{\pi}_{r,j,t}$ is the model-implied pd that corresponds to firms in bin $k_{r,j,t}$. The failure rates $\hat{\pi}_{r,j,t}$ can be computed from filtered or full sample ('smoothed') estimates of the latent risk factors.

Figure 8 reports, for filtered estimates on the left hand side and full-sample based estimates on the right hand side, (a) the time series of \hat{e}_t over time, (b) a QQ plot of the residuals against the normal, (c) an error histogram and associated density kernel estimate, as well as (d) the autocorrelation function (ACF) of errors, respectively. Some larger deviations of predicted from observed values occur around the recession periods of 2001 and 2008-09. Overall, the errors are zero on average and roughly standard normally distributed. The error autocorrelation function suggests that there may be some slight leftover autocorrelation at the first lag, and possibly at the second lag for squared errors. The magnitude of the autocorrelation is not large, however. It may come from imposing relatively strict parameter pooling restrictions for non-financial defaults, which are not the main focus of this paper. When full sample (smoothed) estimates of the default rates are used to predict defaults, most of the larger residuals and error autocorrelation disappear. We conclude that, overall, our model specification and pooling restrictions are an appropriate description of the common dynamics in the default data.