

Systemic risk diagnostics: coincident indicators and early warning signals*

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

^(a) European Central Bank, Financial Markets Research

^(b) Department of Econometrics, VU University Amsterdam

^(c) Tinbergen Institute

^(d) Department of Finance, VU University Amsterdam, and Duisenberg school of finance

August 30, 2011

Abstract

We propose a dynamic factor modeling framework to assess financial system risk. We construct coincident measures (‘thermometers’) and a forward looking indicator (‘barometer’) for the likelihood of a simultaneous failure of many financial intermediaries. The indicators are based on macro-financial and credit risk factors that we extract from a large data set comprising the U.S., the EU-27 area, and the respective rest of the world. Credit risk conditions can significantly and persistently decouple from macroeconomic and financial fundamentals. Our risk indicators may serve as part of a financial sector surveillance and early-warning system for macro-prudential policy.

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

JEL classification: *G21, C33*

*Corresponding author: B. Schwaab, European Central Bank, Kaiserstrasse 29, 60311 Frankfurt, Germany, bernd.schwaab@ecb.int; S.J. Koopman, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, s.j.koopman@feweb.vu.nl; A. Lucas, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, a.lucas@vu.nl. We would like to thank Dion Bongaerts, Christian Brownlees, Michael Dempster, Robert Engle, Philipp Hartmann, Martin Scheicher, and Michel van der Wel for comments. We are further grateful to Thomas Kostka and Silviu Oprica for excellent research assistance at ECB, and seminar participants at Bundesbank, Bank of England, UC Dublin, Erasmus University Rotterdam, European Central Bank, the European Systemic Risk Board secretariat, International Monetary Fund, NYU Stern, SUERF Bruxelles, Tinbergen Institute, VU University Amsterdam, and ZEW Mannheim. 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.

“One of the greatest challenges ... at this time is to restore financial and economic stability. ... The academic research community can make a significant contribution in supporting policy-makers to meet these challenges. It can help to improve analytical frameworks for the early identification and assessment of systemic risks.” Jean-Claude Trichet, President of the ECB, Clare Distinguished Lecture in Economics and Public Policy, University of Cambridge, December 2009.

1 Introduction

Macro-prudential oversight seeks to focus on safeguarding the financial system as a whole. This has proven to be a major issue in the wake of the recent financial crisis. The debate on macro-prudential policies and potential warning signals ignited by the crisis is currently under full swing. Many of the models constructed before the crisis have fallen short in this respect. For example, regulators have learned the hard way that cross-sectional correlations between asset and credit exposures can have severe consequences, even though each of these exposures might be qualified as safe when considered in isolation. Cross-sectional dependence undermines the benefits of diversification and may lead to a ‘fallacy of composition’ at the systemic level, see for example Brunnermeier, Crocket, Goodhart, Persaud, and Shin (2009). In particular, traditional risk-based capital regulation at the individual institution level may significantly underestimate systemic risk by neglecting the macro impact of a joint reaction of financial intermediaries to a common shock.

There is widespread agreement that financial systemic risk is characterized by both cross-sectional and time-related dimensions; see, for example, Hartmann, de Bandt, and Alcalde (2009). The cross-sectional dimension concerns how risks are correlated across financial institutions at a given point in time due to, for example, direct and indirect linkages across institutions and prevailing default conditions. The time series dimension concerns the evolution of systemic risk over time due to, for example, changes in the default cycle, changes in financial market conditions, and the potential buildup of financial imbalances such as asset and credit market bubbles.

In contrast to a broader consensus on the set of models, indicators, and analytical tools for macroeconomic and monetary policy analysis, such agreement appears to be absent for macro-prudential policy analysis. Therefore, the paper aims to make two contributions on systemic risk assessment. First, we develop a unified econometric framework for the measurement of global macro-financial and credit risk conditions. The framework is based on the mixed-measurement dynamic factor model (MM-DFM) approach that is introduced by Koopman, Lucas, and Schwaab (2010). Our model provides a diagnostic tool that tracks

the evolution of macro-financial developments and point-in-time credit risk conditions, as well as their joint impact on the likelihood of failure of a large number of financial intermediaries. Such joint failures are akin to financial earthquakes or tsunamis - low probability events during most times, but with an asymmetrically large and potentially devastating impact on the real economy should the risk materialize. Second, we develop a set of coincident and forward looking indicators for financial distress based on the empirical output of our analysis. We distinguish ‘thermometers’ from ‘barometers’. Thermometers are coincident risk indicators that, metaphorically, a policy maker can plug into the financial system to read off its current ‘heat’. In contrast, a barometer measures current (air) pressure and contains information about (weather) conditions at a later point in time. In a financial stability context, early warning ‘barometers’ may signal whether risks to financial stability, such as an asset price or lending bubble, are currently building up. We suggest that an early warning indicator may be based on deviations of credit risk conditions over time from a set of macro-financial fundamentals, and show that such deviations have preceded macro-financial distress in the past.

We use our modeling framework to study systemic risk conditions across three broad geographical regions, i.e., (i) the U.S., (ii) current EU-27 countries, and (iii) all remaining countries. In this way, our perspective departs substantially from other studies that typically focus on one region only, in particular the U.S. Several people have stressed the importance of such an international perspective, see e.g. de Larosiere (2009), and Brunnermeier et al. (2009). It requires one to look beyond domestic developments for detecting financial stability risk. In the context of the recent crisis. For example, 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, developments in the U.S. housing market have triggered distress for European financial institutions. In the MM-DFM model, we allow for the differential impact of world business cycle conditions on regional default rates, spillovers of such macro effects across regions, unobserved regional risk factors, as well as world-wide financial industry sector dynamics.

Our empirical study is based on worldwide credit data for more than 12.000 firms. We differentiate between the impact of macro and financial market conditions on defaults versus autonomous default dynamics, and industry effects. We refer to the autonomous default dynamics as frailty effects, see also Duffie, Eckner, Horel, and Saita (2009). Our empirical findings show that the magnitude of frailty effects can serve as a warning signal for macro-prudential policy makers. Latent residual effects are highest when aggregate default conditions (the ‘default cycle’) diverge significantly from what is implied by aggregate macroeconomic conditions (the ‘business cycle’), for example due to unobserved shifts in

credit supply and lending standards. 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 too benign compared to observed macro and financial data. In either case, a macro-prudential policy maker should be aware of a possible decoupling of systematic default risk conditions from their macro-financial fundamentals. We demonstrate that the mixed measurement dynamic factor model framework provides the means for a timely detection of this decoupling.

Our work is related to two different research directions in the literature. First, 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 distress 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, Duffie, Kapadia, and Saita (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, Lucas, and Monteiro (2008), Koopman and Lucas (2008), Lando and Nielsen (2009), and Duffie et al. (2009), and Azizpour, Giesecke, and Schwenkler (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. In addition, it allows for both macro, frailty, and industry effects. Finally, it provides a unified framework to integrate systemic risk signals from different sources, whether macroeconomic and financial market conditions, equity markets and balance sheet information (via expected default frequencies, EDFs), or actual defaults.

Another research interest relates to our second contribution: the construction of systemic risk measures. 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, 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, distinguishes between global and regional factors, and looks beyond coincident indicators of risk.

The remainder of this paper is set up as follows. In Section 2, we briefly discuss the desirable properties of a good systemic risk measure in the time dimension. Section 3 discusses our econometric framework that is based on a mixed-measurement dynamic factor model. Some details of parameter and factor estimation are given as well. Section 4 presents the data. Section 5 discusses the main empirical results and presents coincident and forward-looking measures of financial distress. Section 6 concludes.

2 What is needed for measuring systemic risk?

In the literature, systemic risk is understood in two different but related ways. First, the ‘systemic risk contribution’ associated with a financial institution corresponds to a negative externality that its failure would have on other firms and the economy at large. It is the extent to which an individual firm ‘pollutes the public good’ of financial stability. Cross sectional rankings of the risk contribution of financial firms can be obtained, see *inter alia* Adrian and Brunnermeier (2009), Acharya, Pedersen, Philippon, and Richardson (2010), Brownlees and Engle (2010), and Huang, Zhou, and Zhu (2010). Such firm specific risk measures are usually subadditive. Therefore, cross-sectional aggregation may not result in a meaningful measure of risk over time. In addition, such aggregated measures may not perform well in risk assessment practice, see e.g. Moreno and Pena (2010). Conversely, however, systemic risk is often understood as the risk of experiencing a systemic event. We follow this second convention. This notion is analogous to directly assessing the total size of the (risk) pie rather than its composition. In this paper, we define systemic risk operationally as the time varying probability of experiencing a simultaneous failure of a large number of currently active financial intermediaries. Such joint failures have in the past turned out to be very costly in real terms, see for example Reinhart and Rogoff (2009, Chapter 10).

We now identify five core features for appropriate indicators of systemic risk in the time dimension. We refer to these features in the next sections where we discuss our econometric framework.

A broader definition of systemic risk: Current tools for financial risk measurement rely on relatively narrow definitions of a systemic event. A more comprehensive framework could be based on e.g. the theoretical work of Goodhart, Sunirand, and Tsomocos (2006) who argue that systemic risk arises from (i) spillover dynamics at the financial industry level, (ii) shocks to the macroeconomic and financial markets environment, and implicitly (iii) the

potential unraveling of widespread financial imbalances. These sources of risk act on observed data simultaneously, and should therefore all be part of a diagnostic framework. Otherwise, incorrect risk attributions may arise. For example, allowing for interconnectedness through business links but not for shared exposure to common risk factors may spuriously attribute dependence to links that do not exist.

International or inter-regional focus: Several studies have stressed the importance of an international perspective, see e.g. Brunnermeier et al. (2009), de Larosiere (2009) and Volcker et al. (2009). As argued in the introduction, an exclusive focus on domestic conditions is inefficient at best and most likely severely misleading. Consequently, a diagnostic tool for systemic risk should incorporate information from various regions and industries.

Macroeconomic/financial conditions: The main source of risk in the banking book is default clustering. Adverse changes in macroeconomic and financial conditions affect the solvency of all, financial and non-financial, firms in the economy. Observed macro-financial risk factors are therefore systematic and a source of cross-sectional dependence between defaults. The resulting default clusters have a first-order impact on intermediaries' profitability and solvency, and therefore on financial stability. As a result, proxies for time-varying macro-financial and credit risk conditions should be at the core of a systemic risk assessment exercise.

Expected default frequencies: Financial institutions rarely default. This is particularly the case in Europe, where we count 12 financial defaults in the period from 1984Q1 to 2010Q2. Data scarcity poses obvious problems for the modeling of shared financial distress and financial default dependence. As a consequence, models based on actual default experience may only give a partial picture of current stress. Other measures of credit risk can complement historical default information. Such information can be obtained from asset markets (equities, bonds, credit default swaps) and possibly be augmented with accounting data. One candidate that integrates information from accounting data (via debt levels) and forward-looking equity markets (via prices and volatilities) are expected default frequencies (EDF) which are based on structural models for credit risk. We include this measure in our empirical analysis. Other information can be added in the form of credit default swaps (CDS) spreads. However, the short length of time series of liquid CDS for individual firms is typically a problem.

Unobserved factors: The time-varying probability of a systemic event is an inherently unobserved processes. Its main drivers are also unobserved: contagion risk at the financial sector level, changes in shared macro-financial conditions, and financial imbalances such as unobserved large shifts in credit supply. Many of these unobserved conditions, however, can be inferred (reverse-engineered) from different sets of observed data. The appropriate

econometric tools for extracting unobserved factors from observed data are collectively known as state space methods.

3 The diagnostic framework

3.1 Mixed-measurement dynamic factor models

We use the mixed-measurement dynamic factor model (MM-DFM) approach as introduced in Koopman, Lucas, and Schwaab (2010). The approach is based on a state space framework and incorporates all desired features as stated in Section 2. The main idea is to estimate the composite factors of unobserved systemic risk using a panel of time series observations. Once the unobserved (or latent) risk factors are estimated, we can construct an accurate coincident and forward looking measures of systemic risk.

Credit risk is the main risk in the banking book and time-varying credit conditions are therefore central to systemic risk assessment. Our data sources for assessing credit risk consist of N macroeconomic and financial market variables x_t , default counts y_t obtained from historical information across R regions, and expected default frequencies (EDFs) z_t for S_r financial firms in the r th region for $r = 1, \dots, R$ and for time index $t = 1, \dots, T$. The data is denoted by

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

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

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

for $t = 1, \dots, T$, where x_{nt} represents the value of the n th macroeconomic variable at time period t , $y_{r,jt}$ is the number of defaults for economic region r , cross-section j and time period t , and $z_{r,st}$ is the EDF in economic region r of financial s in time period t , for $n = 1, \dots, N$, $t = 1, \dots, T$, $r = 1, \dots, R$, $j = 1, \dots, J$ and $s = 1, \dots, S_r$. Cross-section j can represent different categories of firms. For example, j can represent industry sector, rating category, firm age cohort, or a combination of these. We assume that all variables x_t , y_t , and z_t are driven by a vector of common dynamic factors, that is f_t . However, our panel data may be unbalanced, such that all variables may not be observed at all time periods.

The model combines normally and non-normally distributed variables. We adopt a standard conditional independence assumption: conditional on latent factors f_t , the measurements (x_t, y_t, z_t) are independent over time and within the cross-section. In our specific case and conditional on f_t , we assume that the elements of x_t are normally distributed with their means as functions of f_t . The default counts $y_{r,jt}$ have a binomial distribution with $k_{r,jt}$

trials and with a probability $\pi_{r,jt}$ that is a function of f_t . The number of trials $k_{r,jt}$ refers to the number of firms and $\pi_{r,jt}$ is the probability of default for a specific cross-section j in region r at time t . The EDFs z_t are transformed to represent a frequency for a quarterly horizon. The corresponding log-odds ratio is defined as $\bar{z}_{r,st} = \log(z_{r,st}/(1 - z_{r,st}))$. We effectively model the log-odds as being a normal variable (conditional on f_t). The factor structure distinguishes macro, regional frailty, and industry-specific effects, denoted by f_t^m , f_t^d , f_t^i , respectively. We therefore have $f_t' = (f_t^{m'}, f_t^{d'}, f_t^{i'})$. The latent factors are the main input for our systemic risk measures which we discuss below.

In the factor model structure we assume that the macroeconomic and financial variables in x_t are only determined by the macro factors while the other observed variables in y_t and \bar{z}_t are determined by all factors,

$$x_{nt}|f_t^m \sim \text{Gaussian}(\mu_{nt}, \sigma_n^2), \quad (4)$$

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

$$\bar{z}_{r,st}|f_t^m, f_t^d, f_t^i \sim \text{Gaussian}(\bar{\mu}_{st}, \bar{\sigma}_s^2). \quad (6)$$

where the means μ_{nt} and $\bar{\mu}_{st}$, and probability $\pi_{r,jt}$ are functions of f_t and where the variances σ_n^2 and $\bar{\sigma}_s^2$ are treated as unknown coefficients. The number of firms at risk $k_{r,jt}$ is known since it is observed from the dataset. The factors in f_t^m capture shared business cycle dynamics in both macro and credit risk data, and are therefore common to x_t , y_t , and \bar{z}_t . The frailty factors in f_t^d are region-specific; they only load on the realized defaults, $y_{r,t}$, and the log-odds of EDFs, $\bar{z}_{s,t}$, from a given region. The frailty and industry factors are independent of observed macroeconomic and financial data. They capture variation due to default risk, above and beyond what is already implied by the macro factors f_t^m . The latent factors in f_t^i affect firms in the same industry. Such factors may arise as a result of default dependence through up- and downstream business links, and may capture the industry-specific propagation of aggregate shocks. Both f_t^d and f_t^i help capture a deviation of default activity from what is implied by macro-financial fundamentals as summarized by f_t^m .

The point-in-time default probabilities $\pi_{r,jt}$ in (5) vary over time due to the shared exposure to the underlying risk factors in x_t , as summarized by f_t^m , to the frailty effects f_t^d , and to the latent industry specific effects f_t^i . We model $\pi_{r,jt}$ as the logistic transform of an index function $\theta_{r,jt}$,

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

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

The panel data dynamics in (1) to (3) are captured by time-varying parameters or unobserved signals which are modeled as functions of the dynamic factors in f_t . In particular, we have

$$\mu_{nt} = c_n + \beta'_n f_t^m, \quad (8)$$

$$\theta_{r,jt} = \lambda_{r,j} + \beta'_{r,j} f_t^m + \gamma'_{r,j} f_t^d + \delta'_{r,j} f_t^i, \quad (9)$$

$$\bar{\mu}_{r,st} = \bar{c}_{r,s} + \bar{\beta}'_{r,s} f_t^m + \bar{\gamma}'_{r,s} f_t^d + \bar{\delta}'_{r,s} f_t^i, \quad (10)$$

where $\lambda_{r,j}$, c_n , and $\bar{c}_{r,s}$ are fixed effects, and risk factor sensitivities β , γ , and δ refer to the loadings on macro factors, frailty factors, and industry-specific factors, 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,

$$\bar{\chi}_{r,j} = \chi_0 + \chi_{1,d_j} + \chi_{2,s_j} + \chi_{3,r_j}, \quad \text{for } \bar{\chi} = \lambda, \beta, \gamma, \delta, \bar{\beta}, \bar{\gamma}, \bar{\delta} \quad (11)$$

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, and $\chi_{3,r}$ is the deviation related to regional effects. Since we assume that the baseline effect χ_0 is nonzero, some of the other coefficients need to be subject to zero constraints to ensure identification. The specification in (11) is parsimonious yet sufficiently flexible to accommodate heterogeneity across regions and industries.

The latent factors are stacked into the vector $f_t = (f_t^{m'}, f_t^{d'}, f_t^{i'})'$. We assume that the elements of f_t follow independent autoregressive dynamics. In our study, we have

$$f_t = \Phi f_{t-1} + \eta_t, \quad \eta_t \sim \text{NID}(0, \Sigma_\eta), \quad (12)$$

where the coefficient matrix Φ and covariance matrix Σ_η are assumed diagonal. Extensions to more complex dynamic structures are straightforward exercises. The autoregressive structure in (12), however, already allows sufficient stickiness in the components of f_t . For example, it allows the macroeconomic factors f_t^m to evolve slowly over time and to capture business cycle dynamics in macro and default data. Similarly, the credit climate and industry default conditions are modeled as persistent processes for f_t^d and f_t^i , respectively. The $m \times 1$ disturbance vector η_t is serially uncorrelated. To ensure the identification of the factor loadings, we impose $\Sigma_\eta = I - \Phi\Phi'$. It implies that $E[f_t] = 0$, $\text{Var}[f_t] = I$, and $\text{Cov}[f_t, f_{t-h}] = \Phi^h$, for $h = 1, 2, \dots$. As a result, the loading coefficients $\beta_{r,j}$, $\gamma_{r,j}$, and $\delta_{r,j}$ in (9) can be interpreted as risk factor volatilities (standard deviations) for the firms in cross section (r, j) . It also leads us to the initial condition $f_1 \sim N(0, \Sigma_0)$ and completes the specification of the factor process.

3.2 Parameter and risk factor estimation

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. The model relies on a parameter vector that contains the coefficients in Φ , λ , β , γ , δ , $\bar{\beta}$, $\bar{\gamma}$, and $\bar{\delta}$. This parameter vector is estimated by the method of simulated maximum likelihood. Since our dynamic factor model partly relies on the binomial density, the likelihood function is not available in a convenient analytical form. We therefore need to evaluate the high-dimensional integral of the likelihood function directly. Numerical integration is not computationally feasible for such high-dimensional cases and, therefore, we rely on Monte Carlo simulation methods for evaluating the likelihood function. As the same random numbers can be used for likelihood evaluations for different parameter vectors, the likelihood is a smooth function of the parameter vector. Hence we can maximize the Monte Carlo likelihood function directly by means of a numerical optimization method. We refer to the Appendix A1 for details on our simulation based estimation procedure for mixed measurement data.

An advantage of using state space methods is the convenient treatment of missing values in the dataset. Missing values can have a strong presence in the panels (1) to (3). For example, some macroeconomic variables in x_t may not be available at the beginning of the sample. Also, default data $y_{r,jt}$ is not available (missing) if there are no corresponding firms at risk, that is $k_{r,jt} = 0$. We refer to the Appendix A2 for the treatment of the many missing values in our setup. Clearly, state space methods provide a natural framework to account for missing entries in the data without any adjustments to the model.

The cross-sectional dimension in the panels (1) to (3) can become very large. High-dimensional measurements can lead to computational problems for any method of estimation. Jungbacker and Koopman (2008) show that state space methods for dynamic factor models with high-dimensional measurements and a low-dimensional state vector become computationally feasible when we transform the panel dataset to a time series of observation vectors that have the same dimension as the factors. The transformation results are only justified for the linear Gaussian measurement model. However, many importance sampling computations as detailed in Appendix A1 rely on an approximating linear Gaussian measurement equation. Appendix A3 demonstrates that we can adapt the results of Jungbacker and Koopman (2008) to nonlinear models for partly non-Gaussian data. These methods are helpful regarding the feasibility of the analyses in our empirical study.

3.3 Four thermometers and a barometer

Using the mixed measurement model set-up, we can construct indicators of financial distress for a specific region or combination of regions. Being based on (8) to (10), such indicators automatically integrate the effects of macro, frailty, and industry effects. We consider five indicators, four coincident measures (‘thermometers’), and one forward-looking early warning indicator (‘barometer’). Thermometers are designed to display the current ‘heat’ in the financial system. Our early warning indicator captures credit market imbalances that are currently building up and may pose a risk to the system at a later stage. Both thermometers and early warning indicators are essential tools to monitor system risk in a forward looking policy context.

The first thermometer is the model-implied financial sector failure rate. The time-varying default probability $\pi_{r,jt}$ in (7) can be interpreted as the fraction of financial intermediaries that are expected to fail over the next three months. We estimate this quantity by aggregating implied rates from the bottom up across banks and financial non-banks. Naturally, high failure rates imply high levels of common financial distress, and thus a higher risk of adverse real economy effects through financial failure.

A second thermometer is the time-varying probability of simultaneous failure of a large number of financial intermediaries, as suggested in Giesecke and Kim (2011). Such intermediaries may be depository institutions, but also insurers, re-insurers, and broker/dealers that provide intermediation services. The latter three categories may be considered part of the parallel banking system. Due to the conditional independence assumption, the joint probability of failure can easily be constructed from the binomial cumulative distribution function and the time-varying financial sector failure rates.

A third thermometer is based on the default signals $\theta_{r,jt}$ in (9). The signals $\theta_{r,jt}$ consist of two terms, $\theta_{r,jt} = [\lambda_{r,j}] + [\beta'_{r,j}f_t^m + \gamma'_{r,j}f_t^d + \delta'_{r,j}f_t^i]$, where the fixed effects $\lambda_{r,j}$ pin down the *through-the-cycle* log-odds of the default rate, and the systematic factors f_t^m , f_t^d , and f_t^i jointly determine the *systematic point-in-time* default conditions. The signals $\theta_{r,jt}$ are Gaussian since all risk factors in f_t are Gaussian. We can therefore standardize these signals to unconditionally standard normally distributed values $z_{r,jt}^\theta$,

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

where $\text{Var}(\theta_{r,jt}) = \beta'_{r,j}\beta_{r,j} + \gamma'_{r,j}\gamma_{r,j} + \delta'_{r,j}\delta_{r,j} \geq 0$ 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

$$\text{SRI}_{r,jt} = \bar{\Phi}(z_{r,jt}^\theta), \quad (13)$$

where $\bar{\Phi}(z)$ is the standard normal cumulative distribution function. Values of $\text{SRI}_{r,jt}$ lie between 0 and 1 by construction with uniform (unconditional) probability. Values below

0.5 indicate less-than-average common default stress, while values above this value suggest above-average stress. Values below 20%, say, are exceptionally benign, and values above 80% are indicative of substantial systematic stress. Other percentiles can similarly be considered. Our measure of financial system risk is obtained when (13) is applied to model-implied failure rates for *financial* firms in a given region.

A fourth indicator of financial system risk is the expected number of financial defaults over the next year conditional of at least one financial default occurring,

$$\text{BSI}_{r,j} = k_{r,jt}\pi_{r,jt} / (1 - \text{Binomial}(0; k_{r,jt}, \pi_{r,jt})). \quad (14)$$

This Banking Stability Index has been proposed by Huang (1992), and subsequently used by e.g. Hartmann et al. (2005) and Segoviano and Goodhart (2009). Naturally, a high expected number of financial defaults indicates adverse financial conditions.

Finally, the indicator (13) can be modified to only capture frailty and industry effects. This yields 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 . This indicator is our ‘credit risk deviations’ early warning indicator,

$$\text{CBI}_{r,j} = (\gamma'_{r,j}f_t^d + \delta'_{r,j}f_t^i) / \sqrt{\gamma'_{r,j}\gamma_{r,j} + \delta'_{r,j}\delta_{r,j}}. \quad (15)$$

Section 5 below reports and discusses the indicator values from this section. 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.

4 Data

We use data from three main sources in the empirical study below. 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. Such macroeconomic and financial market data is usually stressed in a stress testing exercise, and considered for the U.S. and Europe. Table 1 provides a listing. The macro variables enter the analysis as annual growth rates from 1984Q1 to 2010Q4.

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 1984Q1 to 2010Q4. From this data, we construct quarterly values for $y_{r,jt}$ and $k_{r,jt}$ in (5). When counting exposures $k_{r,jt}$ and corresponding defaults $y_{r,jt}$, a previous rating withdrawal is ignored if it is followed by a later default. If there are multiple defaults per firm, we consider

Table 1: International macroeconomic time series data

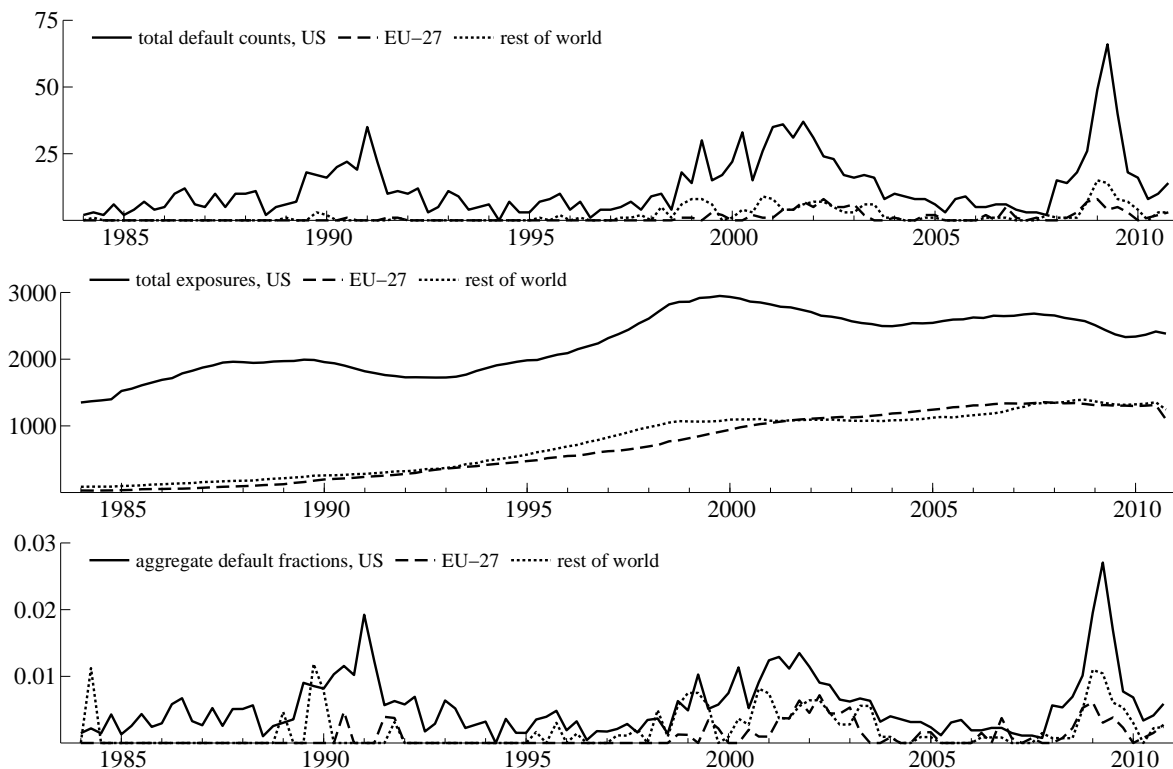
We list the variables contained in the macroeconomic panel. The time series data enters the analysis as yearly (yoy) growth rates. The sample is from 1984Q1 to 2010Q4.

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 U.S. Treasury Bond Yield, 20 years U.S. 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 (Euribor), 3 months Euro Area (EA16) Industrial Confidence Indicator	8

16

Figure 1: Actual default experience

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 different economic regions: the U.S., the EU-27 area, and the respective rest of the world.



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. We use the industry specification to distinguish between financial and non-financial firms.

Table 2 provides an overview of the international exposure and default count data. Corporate data is most abundant for the U.S., with E.U. countries second. Most firms are either from the industrial or financial sector. The bottom of Table 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. Figure 1 plots aggregate default counts, exposures, and observed fractions over time for each economic region.

Table 2 reveals that financial intermediaries rarely default, in particular in Europe. This is an obvious problem for inference on time-varying risk conditions. For financials, we therefore

Table 2: International default and exposure counts

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

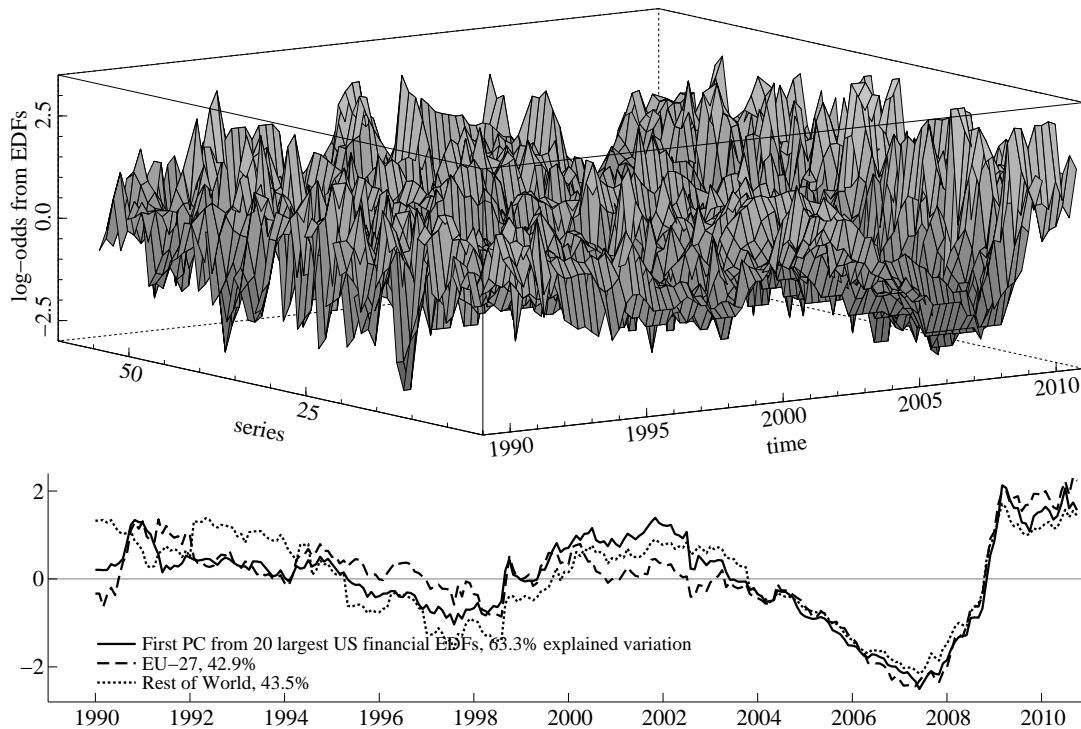
No Defaults	U.S.	Europe	Other	Sum
Bank	41	8	22	71
Fin non-Bank	84	4	14	102
Transport	90	17	8	115
Media	127	2	2	131
Leisure	97	9	15	121
Utilities	24	2	5	31
Energy	79	0	7	86
Industrial	435	16	53	504
Technology	177	38	24	239
Retail	94	1	4	99
Cons Goods	120	8	17	145
Misc	31	0	16	47
Sum	1399	105	187	1691

Firms	U.S.	Europe	Other	Sum
Bank	478	603	591	1672
Fin non-Bank	966	371	500	1837
Transport	336	70	72	478
Media	460	33	34	527
Leisure	434	73	59	566
Utilities	597	149	138	884
Energy	512	84	152	748
Industrial	1920	419	497	2836
Technology	941	204	220	1365
Retail	311	32	46	389
Cons Goods	591	110	112	813
Misc	250	151	258	659
Sum	7796	2299	2679	12774

Firms, 2008Q1	U.S.	Europe	Other	Sum
Aaa	50	84	69	203
Aa	141	355	250	746
A	415	403	337	1155
Baa	575	229	291	1095
Ba	278	72	177	527
B	673	96	183	952
Ca-C	379	58	56	493
Sum	2511	1297	1363	5171

Figure 2: Expected default frequencies of 60 global financials

The top panel reports the standardized log-odds from EDF data for the largest 60 global financial firms (banks and financial non-banks). The sample consists of the largest 20 U.S., EU-27, and rest of the world financial firms, respectively. The raw data sample is from 1990Q1 to 2010Q4, and contains missing values. Missing values are inferred using the EM algorithm of Stock and Watson (2002). The bottom graph plots the respective first principal components from the U.S. , EU-27, and the rest of the world sub-sample.



add data from a third dataset. Data on expected default frequencies for the 20 largest (based on 2008Q4 market cap) financial firms in the U.S. , EU-27, and respective rest of the world, is taken from Moody’s KMV CreditEdge. These $3 \times 20 = 60$ expected default frequencies are based on a firm value model that takes equity values and balance sheet information as inputs. We use it to augment our relatively sparse data on actual defaults for financial firms. Figure 2 plots the panel of EDF data, after transformation to a quarterly scale and log-odds ratio. The principal components and reported eigenvalues in the bottom panel indicate substantial common variation across institutions and regions that can be summarized in a factor structure.

5 Empirical results on system risk

This section presents the main empirical findings. Section 5.1 comments briefly on the main sources of financial default clustering. Sections 5.2 and 5.3 present our thermometers and forward looking indicator for systemic risk assessment.

5.1 Why do financial defaults cluster?

Observed credit risk data reveals that aggregate financial sector failure rates are up to ten times higher in bad times than in good times. This is striking. Why do financial failures cluster so dramatically over time? Which sources of risk are important, and to what extent? The answer to these questions is important for constructing effective coincident and forward looking risk indicators.

Based on preliminary data analysis and testing, relying on results in Bai and Ng (2002) and Alessi, Barigozzi, and Capasso (2010), we report estimation results for a favorite specification with *four* macro factors f_t^m , *three* region and default data specific factors f_t^d , one for each economic region, and one financial industry specific risk factor f_t^i . The latter loads on financial firms from all regions but possibly to different degrees. The four macro-financial factors f_t^m are common to all macro covariates, thus taking into account the positive correlation between U.S. and E.U. conditions that may result from macroeconomic linkages.

Table 3 presents the parameter estimates for model specification (1) to (12). The fixed effects and factor loadings in the signal equation (9) satisfy the additive structure (11). Coefficients λ in the left column combine to the baseline failure rates. The middle and right-hand columns present estimates for loadings β , γ , and δ that pertain to macro, frailty, and industry factors, respectively.

The parameter estimates indicate that macro, frailty, and industry effects are all important for international credit risk conditions. Defaults from all regions and industries load

Table 3: Parameter estimates

We report the maximum likelihood estimates of selected coefficients in the specification of the log-odds ratio (9) with parameterization (11) for λ and β . Coefficients λ combine to fixed effects, or baseline failure rates. Factor loadings β , γ , and δ refer to four 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 vice versa. The frailty factors load on financial and non-financial firms' defaults in a respective region. The estimation sample is from 1984Q1 to 2010Q4.

Intercepts λ_j ,			Loadings f_t^m			Loadings f_t^d, f_t^i		
par	val	t-val	par	val	t-val	par	val	t-val
λ_0	-5.30	11.56	$\beta_{1,0}$	-0.18	2.92	$\gamma_{US,0}$	0.12	1.86
			$\beta_{1,1,fin}$	0.20	2.15	$\gamma_{US,1,fin}$	0.14	1.72
$\lambda_{1,fin}$	-1.09	5.09	$\beta_{1,2,EU}$	0.03	0.46	$\gamma_{EU,0}$	-0.66	3.28
$\lambda_{2,EU}$	-1.64	4.06	$\beta_{1,2,RoW}$	-0.37	3.89	$\gamma_{EU,1,fin}$	0.88	4.06
$\lambda_{2,RoW}$	-1.06	4.02				$\gamma_{RoW,0}$	0.30	1.31
			$\beta_{2,0}$	0.12	2.04	$\gamma_{RoW,1,fin}$	0.61	2.74
			$\beta_{2,1,fin}$	0.10	1.13			
			$\beta_{2,2,EU}$	0.06	1.00	δ_{US}	0.42	5.68
			$\beta_{2,2,RoW}$	-0.15	1.93	δ_{EU}	0.53	6.71
						δ_{RoW}	0.78	6.78
			$\beta_{3,0}$	0.10	1.80			
			$\beta_{3,1,fin}$	-0.12	1.71			
			$\beta_{3,2,EU}$	-0.06	1.00			
			$\beta_{3,2,RoW}$	0.10	1.29			
			$\beta_{4,0}$	0.58	7.64			
			$\beta_{4,1,fin}$	-0.23	2.16			
			$\beta_{4,2,EU}$	-0.07	0.83			
			$\beta_{4,2,RoW}$	0.16	1.48			

Table 4: Why do financial defaults cluster?

We report the results of a variance decomposition of transformed (Gaussian, log-odds) failure rates for financial firms in three economic regions. The unconditional variance is attributed to three latent sources of financial distress. Each source of distress is captured by a corresponding set of latent factors and associated risk factor standard deviations. Specifically, $s_{r,j}^m = \beta'_{r,j}\beta_{r,j}/\text{Var}(\theta_{r,jt})$, $s_{r,j}^d = \gamma'_{r,j}\gamma_{r,j}/\text{Var}(\theta_{r,jt})$, and $s_{r,j}^i = \delta'_{r,j}\delta_{r,j}/\text{Var}(\theta_{r,jt})$, where $\text{Var}(\theta_{r,jt}) = \beta'_{r,j}\beta_{r,j} + \gamma'_{r,j}\gamma_{r,j} + \delta'_{r,j}\delta_{r,j} \geq 0$, and j refers to financial firms. The estimation sample is from 1984Q1 to 2010Q4.

	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.	44.7%	15.9%	42.4%
EU-27	33.2%	9.1%	57.7%
Rest of world	35.9%	9.0%	55.1%

significantly on common factors from global macro-financial data. This by itself already implies a considerable degree of default clustering. In general, however, common variation with macro data is not sufficient. Frailty effects are found to be important in all regions. The financial industry-specific factor also loads significantly on data from all regions. The shared risk dynamics for financial firms across regions are intuitive: The financial failure dynamics are mainly determined by the EDF data for large and complex banking groups; these groups in turn operate globally both in terms of their lending and funding activities.

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, excess default clustering 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 shocks to macroeconomic and financial conditions are an important source of financial distress. Historically, financial sector stress and business cycle downturns have tended to occur at roughly the same time. This is intuitive, since financial stress may have negative real consequences, and vice versa, possibly with significant feedback and amplification effects. Timing effects are only captured indirectly in this decomposition, as current estimates of f_t^m capture a rotated version of current and lagged structural driving forces, see Stock and Watson (2002) for a discussion and intuition from the linear Gaussian context. Table 4 further suggests that financial industry specific dynamics are an important

additional source of joint financial failure. As a result, financial sector risk dynamics can differ substantially from what is implied by shared exposure to observed macro-financial covariates. We conclude that, in general, all three sources of risk should all be accounted for. Also, deviations of credit risk conditions from macro-financial conditions are potentially large and important for financial sector risk in each region.

5.2 Thermometers: coincident indicators of financial distress

This section presents the coincident risk indicators that are constructed from the estimated risk factors and loading parameters. Figure 3 plots a model-implied failure rate for a large cross section of E.U. and U.S. financial firms. The failure rate is the share of overall intermediaries that can be expected to fail over the next three months. The aggregate rates are obtained by aggregating from the bottom up across approximately 450 U.S. and 400 E.U. area financial firms, respectively. As a result, the reported failure rates take into account a significant part of the parallel banking system, i.e., insurers, real estate firms, and other rated non-bank financial firms that play a role in the intermediation process. Figure 3 also compares the aggregate failure rates (solid lines) with the mean expected default probability (from Moody's KMV) from the largest twenty financial firms (according to 2008Q4 market cap; dashed lines) in the U.S. (left panel) and E.U. area (right panel), respectively. The financial distress in the U.S. during the recession years of 1991, 2001, and 2007-10 are visible from the left panel. However, a recession is not necessary for such stress. An example is the period in the late 1980s in the U.S., when common stress is pronounced while the economy is not in recession. Finally, the financial sector failure rate is different and almost always higher than what is suggested by an analysis of the average EDFs for the twenty largest - and highly rated - large and complex banking groups in each region.

Both the mean EDF and the model-implied rate suggest high levels of common stress for U.S. and E.U. financial firms at the end of the sample. For U.S. financials, the quarterly failure rate ranges from slightly above zero in good times to approximately 1% at the peak of the 2008 crisis. As of 2010Q4, U.S. rates are such that approximately 3% of active financial intermediaries of average size can be expected to fail if this level of stress prevailed during 2011. Model-implied stress for European intermediaries is lower than for U.S. financials. This is partly due to sample composition: rated European financials that tap the financial markets tend to have a high credit rating on average, see Table 2, and the observed historical default frequencies are very low.

Systemic risk is necessarily a multivariate concept, involving a system of banks and financial non-banks. Systemic risk is understood as the risk of experiencing a simultaneous failure of a large number of financial institutions. Simultaneous bank failures are analogous

Figure 3: Implied financial sector failure rate

We plot the model-implied default failure rate for financial sector firms. The sector failure rate is obtained by aggregating across banks and financial non-banks from the bottom up. The estimation sample is 1984Q1 to 2010Q4.

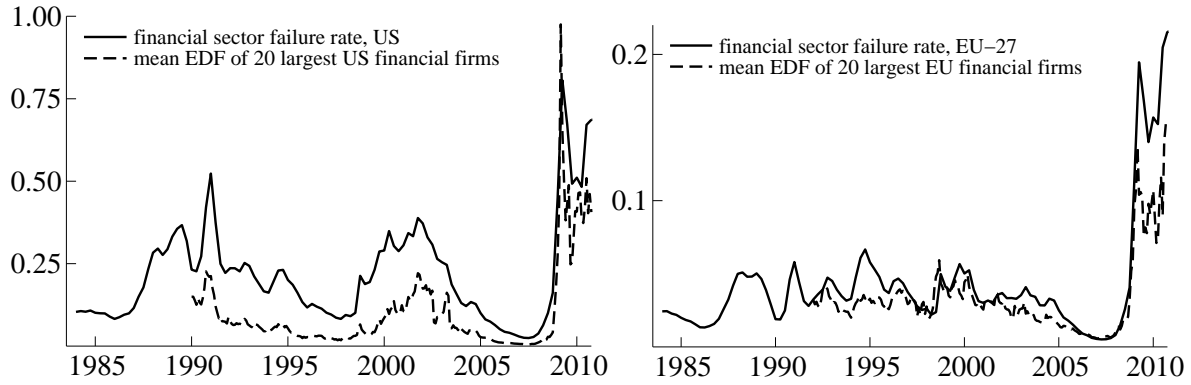


Figure 4: Probability of simultaneous financial failures

We report the probability of a systemic event, defined as the simultaneous failure of $k\%$ or more financial firms over a one year period. The horizontal axis measures time from 1984Q1 to 2010Q4. The vertical axis measures the time-varying probability as a decreasing function of k . The left and right panel refers to the U.S. and the E.U. area, respectively.

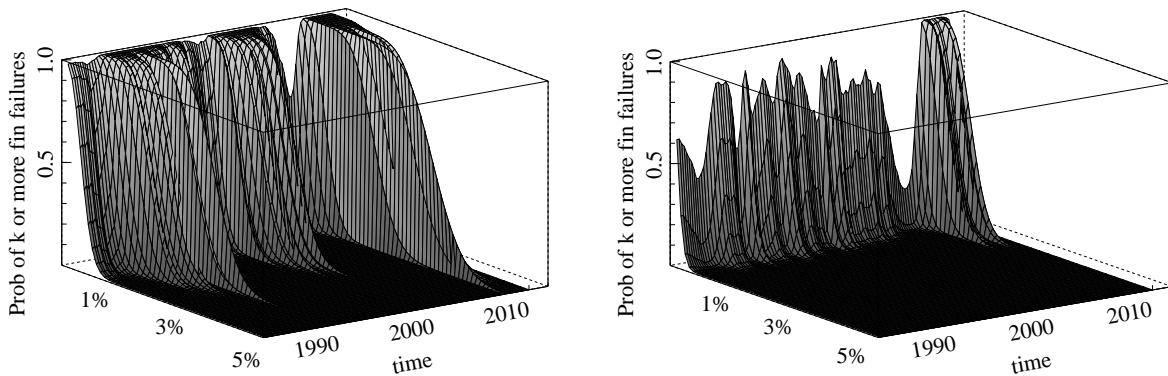
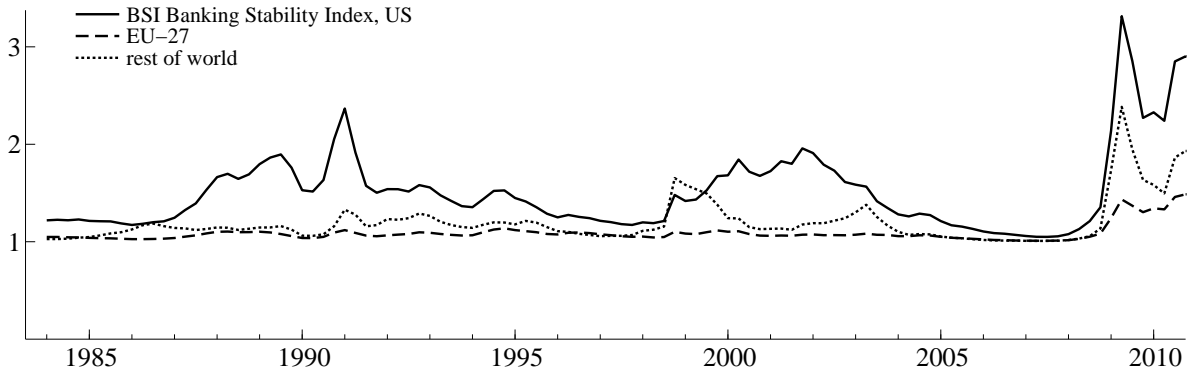


Figure 5: Banking stability index

We plot indicator (14), i.e., the expected number of financial defaults over a one year horizon conditional on at least one default occurring. Firms at risk are held fixed at 100. Financial firms comprise banks and financial non-banks.



to natural disasters such as earthquakes and tsunamis - unlikely events during most times, but with an asymmetrically large and potentially devastating impact if the risk materializes.

The top panels in Figure 4 plot the probability of at least $k\%$ of financial firms failing over a one year horizon (vertical axis), as a decreasing function of k , over time from 1984Q1 to 2010Q4 (horizontal axis). The left and right panels refer to the U.S. and E.U., respectively. The bottom panels cut the three-dimensional plots into various slices along the time dimension, at 0.1%, 0.5%, 1%, and 2% of overall financial sector firms. The figure reveals that, for example, in the E.U. area in 2010Q4, the probability of failure of at least 1% of financial sector firms (e.g., at least four firms of average size out of four hundred firms), at coincident levels of stress, is around 30%. This is a substantial risk of simultaneous failings.

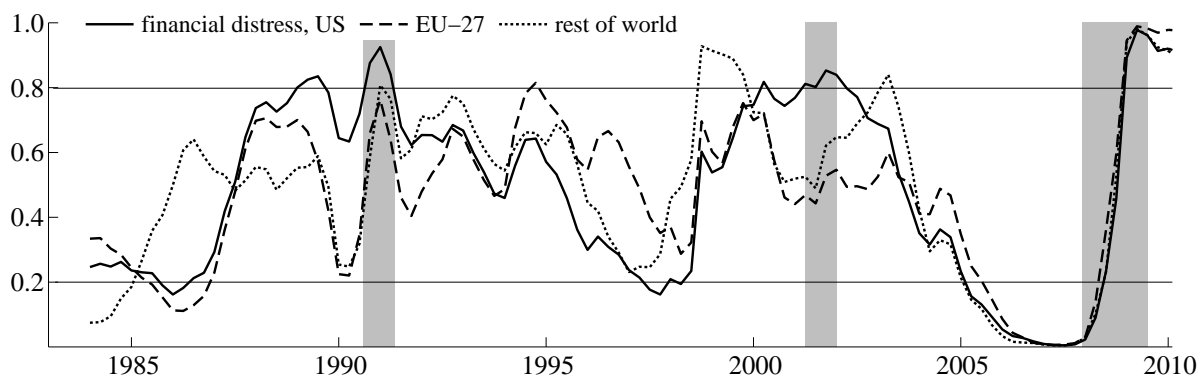
Figure 5 plots the expected number of financial defaults over a one year horizon conditional on at least one firm failing over that time period, see (14). The number of firms is fixed at 100. During the peak of the financial crisis, about three from one hundred U.S. financials are expected to fail over a one year horizon conditional on one firm going down. No data on financial sector counterparty exposures is used for this estimate.¹

Finally, Figure 6 plots financial distress based on the indicator (13). The probability integral transformation in (13) maps common financial distress into a uniform random variable, such that its percentiles can be read off the transformed y-axis. We refer to the best and worst 20% of times as relative ‘exuberance’ and ‘crisis’, respectively, but note that other

¹The newly founded U.S. Office of Financial Research (OFR) is mandated to make an important step in obtaining such data, and has a strong backing through the Dodd Frank act. The OFR may set data standards and has legal subpoena power to obtain information from financial institutions. As of now and the near future, however, interbank counterparty exposures are simply not observed.

Figure 6: Scaled financial distress

The figure plots the risk indicator (13) based on the model-implied financial sector failure rates. A percentile-to-percentile transformation implies that relative levels of implied distress can be read off the y-axis. The best and worst 20% of times are referred to as times of exuberance and crisis, respectively.



thresholds can also be used. Financial distress is virtually absent during the late-1990s and mid-2000s. The late 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. We note that bubbles may have started to build up at each of these times (the dot.com and a lending bubble, respectively). The role of a macro-prudential policy maker may then be to communicate such developments, and to consider taking away the punch bowl from exuberant market participants if a mandate and a careful consideration of the involved tradeoffs allow to do so. Conversely, support measures may be justified during times of crisis. The indicator (13) helps in assessing the relative severity of stress to this purpose.

5.3 An early warning barometer

Past experiences of financial fragility, financial booms and financial crises, suggests that problems rarely appear at the same place in the financial system twice in a row. The main commonality between the different events that turned into a fully fledged financial crisis is that they were not widely expected by market participants and regulators. Goodhart and Persaud (2008), for example, point out that if market prices for assets or credit were good at predicting crashes, crises 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 an asset or

lending bubble. These findings suggest that (i) building early warning signals solely based on market prices has obvious drawbacks, and that (ii) it may be useful to look for structure in the ‘unexpected’, or leftover variation.

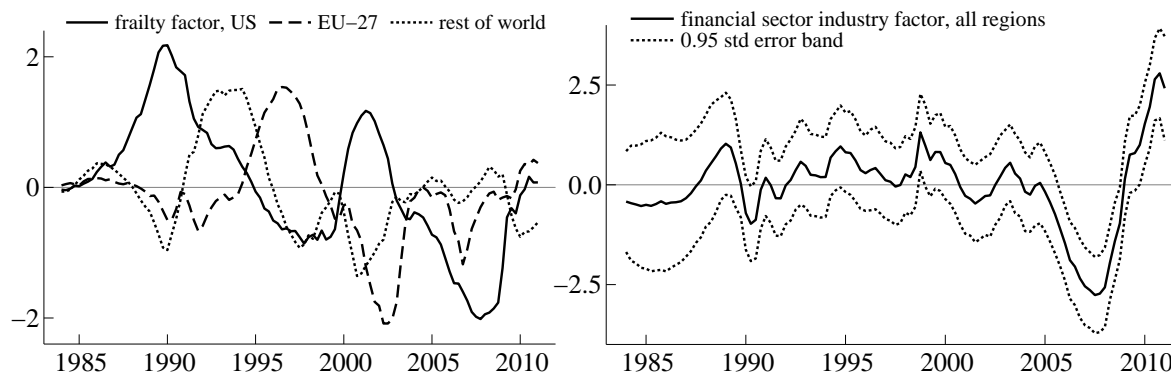
Our warning signals build on Duffie, Eckner, Horel, and Saita (2009), Azizpour, Giesecke, and Schwenkler (2010), and Koopman, Lucas, and Schwaab (2011) who find substantial evidence for a dynamic unobserved risk factor driving default for U.S. firms above and beyond what is implied by observed macro-financial covariates and other information. We interpret the 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). 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 reported in Maddaloni and Peydro (2011). These findings suggest that frailty, among other effects, captures outward shifts in (unobserved) credit supply. Changes in the ease of credit access affect credit risk conditions: it is hard to default if one is drowning in credit. Conversely, even healthy firms may come under stress if credit is rationed at the economy wide level. As a result, systematic default risk (‘the default cycle’) can decouple from what is implied by macro-financial conditions (‘the business cycle’).

Tracking credit risk conditions and deviations from fundamentals is related to the notion of tracking credit aggregates over time. The usefulness of the private credit to GDP-ratio as an early warning indicator for costly asset price busts and banking crisis is a recurring finding in the respective literature, see Borio and Lowe (2002), Misina and Tkacz (2009), Alessi and Detken (2011), Borio and Drehmann (2009), and Barrell, Davis, Karim, and Liadze (2010). The main difference is that we suggest to track credit *risk* instead of credit *quantities* relative to macro-financial conditions.

The left panel of Figure 7 presents the estimated frailty factors for the U.S., EU-27, and the rest of the world. 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. At these times, frailty effects are large in absolute value, and significantly different from zero. On top of these developments, the financial industry-specific factor estimate indicates a particularly benign risk environment for financial firms during the years leading up to the crisis, and

Figure 7: Latent factor estimates

The left panel reports the conditional mean estimates of three region-specific frailty factors. The right panel plots the financial sector industry factor that is common to financial firms in all regions. The approximate standard error bands in the right panel are at a 95% confidence level.



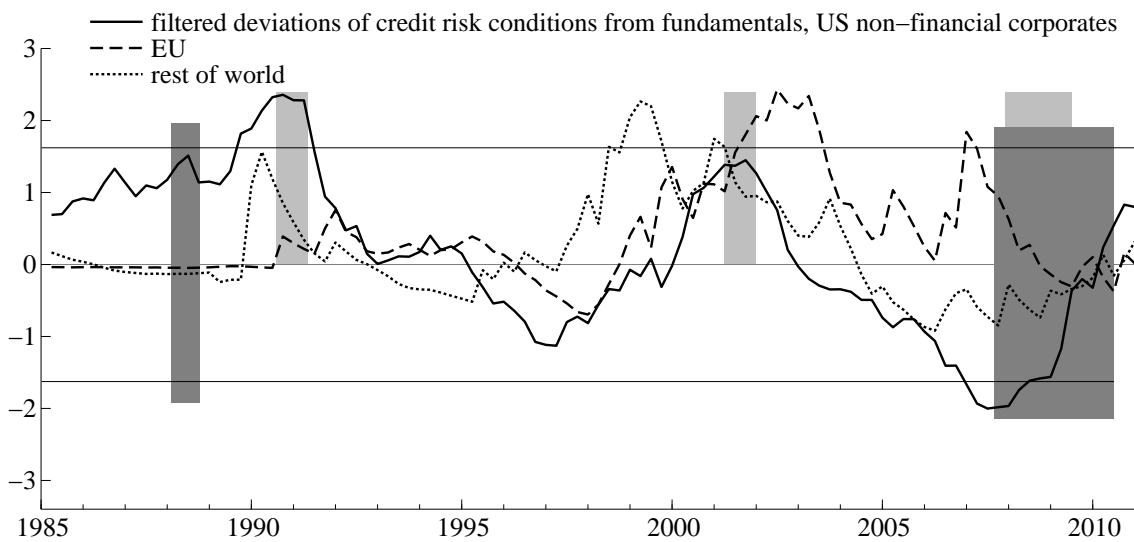
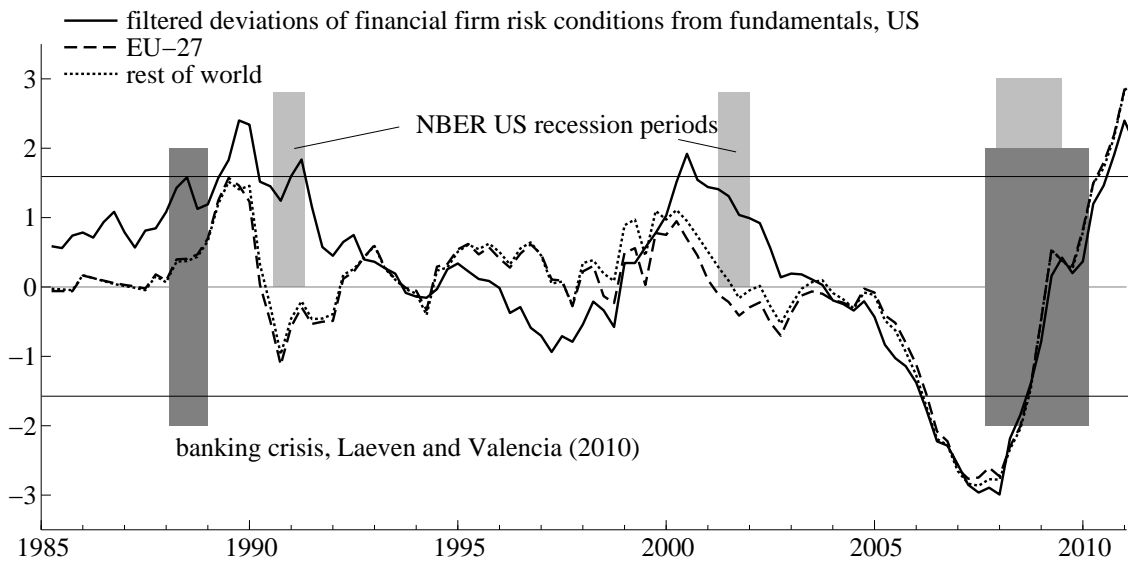
indicates additional stress during and after the 2007-09 crisis.

The top and bottom panels of Figure 8 plot our “credit risk deviations” early warning indicator (15) for financial and non-financial firms, respectively. The indicator combines estimated frailty and industry effects into a warning signal. The indicator captures the extent to which local stress in a given industry sector differs from that which macro-financial fundamentals would suggest. The figure compares estimated deviations in U.S., E.U. area, and respective rest of the world. 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, i.e., risk factor estimation takes into account only information available at that time. Factor loadings are fixed at end of sample values.

The indicator (15) is a standard normal covariate by construction. There are three banking crises in the sample, two in the U.S. (1988 and 2007-10) and one in the E.U. area (2008-10). Figure 8 draws signal thresholds at 1.64 which correspond to a 90% confidence level for a standard normal covariate. Figure 8 reveals that a particularly large and persistent decoupling of risk conditions from fundamentals for both financial and non-financial firms has preceded 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. Cases in point are U.S. conditions during the years 1988-90 leading up to the

Figure 8: ‘Credit risk deviations’ early warning indicator

We plot deviations of credit risk conditions (here for financial firms) from macro-financial fundamentals as captured by the indicator (15). Light and dark grey areas correspond to NBER U.S. recession periods and times of banking crises as identified in Laeven and Valencia (2010), respectively. The indicator is a standard normal variable by construction; the horizontal lines are drawn at one standard deviation.



1991 recession, and conditions during 2010Q3-Q4 in all regions. High financial sector stress may raise the risk of a credit crunch through the bank lending channel, see for example Altunbas, Gambacorta, and Marques-Ibanez (2010). For financials firms, the end of sample developments may reflect in part that the macro fundamentals do not take into account Euro Area sovereign default risk conditions. Such stress is then captured by a latent (industry) factor, see Figure 7.

While a full comparison of alternative indicators for different policy preferences and sets of countries is clearly outside the scope of this paper, we find that our deviations of E.U. filtered risk conditions from fundamentals in excess of their 90th percentile in absolute value do well in identifying the costly asset-price booms in the E.U. area as determined in Alessi and Detken (2011). These asset price booms are costly because they precede macroeconomic and financial distress, and usually coincide with sizable expansions in credit. For Alessi and Detken's data on boom and busts in 11 E.U. countries from 1985Q1 to 2008Q4, and our corresponding EU-27 indicator as graphed in Figure 8, we find a noise to signal ratio of 0.23, an average lead time of 3.89 quarters preceding the bust, and a 70% conditional probability of being in a costly boom if a signal occurs. These are encouraging results. We also confirm the 90% optimal threshold for a broad set of policy preferences.

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. 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 related but inherently different processes. The extent to which they have decoupled is indicated by our early warning barometer.

6 Conclusion

We proposed a novel diagnostic framework for financial systemic risk assessment based on a mixed-measurement dynamic factor model. We combined the risk factor and parameter estimates into new and straightforward coincident and forward looking indicators of financial system risk. Conceptually, our factor structure allows us to address the computational challenges associated with a large cross-sectional dimension of firms more easily than alternative frameworks for financial stability assessments. The new method easily allows one to combine different sets of panel data in a single integrated framework. In our empirical analysis, we found that a decoupling of credit risk from macro-financial fundamentals may serve as an

early warning signal for a macro-prudential policy maker.

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

$$p(y; \psi) = \int p(y, f; \psi) df = \int p(y|f; \psi) p(f; \psi) df, \quad (\text{A.16})$$

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

$$p(y; \psi) = \int \frac{p(y, f; \psi)}{g(f|y; \psi)} g(f|y; \psi) df = g(y; \psi) \int \frac{p(y|f; \psi)}{g(y|f; \psi)} g(f|y; \psi) df. \quad (\text{A.17})$$

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

$$\tilde{p}(y; \psi) = g(y; \psi) M^{-1} \sum_{k=1}^M \frac{p(y|f^{(k)}; \psi)}{g(y|f^{(k)}; \psi)}, \quad f^{(k)} \sim g(f|y; \psi), \quad (\text{A.18})$$

is numerically 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

$$y_{jt} = \mu_{jt} + \theta_{jt} + \varepsilon_{jt}, \quad \varepsilon_{jt} \sim N(0, \sigma_{jt}^2), \quad (\text{A.19})$$

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

To simulate values from the importance density $g(f|y; \psi)$, the simulation smoothing method of Durbin and Koopman (2002) can be applied to the approximating model (A.19). For a set of M draws of $g(f|y; \psi)$, the evaluation of (A.18) 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 $y_{jt} - \mu_{jt} - \theta_{jt} \sim N(0, \sigma_{jt}^2)$, that is (A.19), and $g(y; \psi)$ can be computed by the Kalman filter applied to (A.19), 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.18) 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: treatment of missing values

When missing values are present in the data vector $y = (y'_1, \dots, y'_T)'$, some care must be taken when computing the importance sample weights $w_k = p(y|f^{(k)}; \psi)/g(y|f^{(k)}; \psi)$, $f^{(k)} \sim g(f|y; \psi)$. The mode estimates of the corresponding signals $\theta = (\theta'_1, \dots, \theta'_T)'$ and factors $f = (f'_1, \dots, f'_T)'$ are available even when we have missings. Some bookkeeping is required to evaluate $p(y|f; \psi)$ and $g(\tilde{y}|f; \psi)$ at the corresponding values of f , or θ . Forecasts \tilde{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 by applying the estimation and signal extraction techniques of Section 6 to data (y_0, \dots, y_{T+H}) .

Appendix A3: collapsing observations

Jungbacker and Koopman (2008) show that an $[N \times 1]$ vector of (Gaussian) observations y_t can be collapsed into an $[m \times 1]$ vector of transformed observations \tilde{y}_t^l with $m < N$ usually much smaller than N . The transformation does not lead to loss of information for the estimation of the factors f_t via the Kalman filter and smoother (KFS). We show how this argument can be used in a nonlinear mixed-measurement setting. We focus on collapsing the artificial Gaussian data \tilde{y}_t with associated covariance matrices \tilde{H}_t , see (A.19) and (12).

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 \tilde{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 : 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 $C_t' C_t = (Z' \tilde{H}_t^{-1} Z)^{-1}$ 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 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 .

We refer to Jungbacker and Koopman (2008, Illustration 4) for further details.

References

- Abreu, D. and M. K. Brunnermeier (2003). Bubbles and crashes. *Econometrica* 71(1), 173–204.
- 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., M. Barigozzi, and M. Capasso (2010). Improved penalization for determining the number of factors in approximate factor models. *Statistics & Probability Letters* (80), 1806–1813.
- 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*, forthcoming.
- Altunbas, Y., L. Gambacorta, and D. Marques-Ibanez (2010). Bank risk and monetary policy. *Journal of Financial Stability* 6(3), 121–129.
- Azizpour, S., K. Giesecke, and G. Schwenkler (2010). Exploring the sources of default clustering. *Stanford University working paper series*.
- Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica* 70(1), 191–221.
- Barrell, R., E. P. Davis, D. Karim, and I. Liadze (2010). Bank regulation, property prices and early warning systems for banking crises in oecd countries. *Journal of Banking and Finance* 34, 2255–2264.
- Borio, C. and M. Drehmann (2009). Assessing the risk of banking crises - revisited. *BIS Quarterly Review*, March 2009.
- Borio, C. E. V. and P. W. Lowe (2002). Asset prices, financial and monetary stability: Exploring the nexus. *BIS working paper* 114(7).
- Brownlees, C. T. and R. Engle (2010). Volatility, correlation, and tails for systemic risk measurement. *NYU working paper*.

- 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 2007/2008. *Journal of Economic Perspectives* 23(1), 77–100.
- 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 Larosiere, J. (2009). *The high level group on financial supervision in the E.U. - Report*. Bruxelles: The European Commission.
- Duffie, D., A. Eckner, G. Horel, and L. Saita (2009). Frailty Correlated Default. *Journal of Finance* 64(5), 2089–2123.
- Durbin, J. and S. J. Koopman (1997). Monte Carlo Maximum Likelihood estimation for non-Gaussian State Space Models. *Biometrika* 84(3), 669–684.
- Durbin, J. and S. J. Koopman (2001). *Time Series Analysis by State Space Methods*. Oxford: Oxford University Press.
- Durbin, J. and S. J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika* 89(3), 603–616.
- 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, forthcoming*.
- Goodhart, C. and A. Persaud (2008). A proposal for how to avoid the next crash. *Financial Times UK, 31 Jan 2008*.
- Goodhart, C. A. E., P. Sunirand, and D. P. Tsomocos (2006). A model to analyse financial fragility. *Annals of Finance* 2(1), 1–21.
- Hartmann, P., O. de Bandt, and J. L. P. Alcalde (2009). Systemic risk in banking: An update. *Oxford Handbook of Banking, ed. by A. Berger, P. Molyneux and J. Wilson, Oxford University Press*.
- Hartmann, P., S. Straetmans, and C. de Vries (2005). Banking system stability: A cross-atlantic perspective. *NBER working paper 11698*, 1–86.
- Huang, X., H. Zhou, and H. Zhu (2010). Systemic risk contributions. *Working paper*.
- Jungbacker, B. and S. J. Koopman (2008). Likelihood-based Analysis of Dynamic Factor Models. *Tinbergen Institute Discussion Paper 08-007/4*.

- 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. 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 (2010). Macro, frailty, and contagion effects in defaults: Lessons from the 2008 credit crisis. *Tinbergen Institute Discussion Paper 2010-004/2*, 1–40.
- Koopman, S. J., A. Lucas, and B. Schwaab (2011). Modeling frailty correlated defaults using many macroeconomic covariates. *Journal of Econometrics* 162, 312–325.
- 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.
- Misina, M. and G. Tkacz (2009). Credit, asset prices, and financial stress in Canada. *International Journal of Central Banking* (4), 95–122.
- Moreno, M. R. and J. I. Pena (2010). Systemic risk measures: the simpler the better? . *working paper*.
- Reinhart, C. M. and K. S. Rogoff (2009). *This time is different: Eight centuries of financial folly*. Princeton and London: Princeton University Press.
- 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 (2002). Forecasting Using Principal Components From a Large Number of Predictors. *Journal of the American Statistical Association* 97(460), 1167–1179.
- Trichet, J.-C. (2009). Clare Distinguished Lecture in Economics and Public Policy. *University of Cambridge, 10 December 2009*.
- 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.