

Financial Network Systemic Risk Contributions*

NIKOLAUS HAUTSCH¹, JULIA SCHAUMBURG² and
MELANIE SCHIENLE³

¹University of Vienna; ²VU University Amsterdam and ³Leibniz University
Hannover

Abstract. We propose the realized systemic risk beta as a measure of financial companies' contribution to systemic risk, given network interdependence between firms' tail risk exposures. Conditional on statistically pre-identified network spillover effects and market and balance sheet information, we define the realized systemic risk beta as the total time-varying marginal effect of a firm's Value-at-risk (VaR) on the system's VaR. Statistical inference reveals a multitude of relevant risk spillover channels and determines companies' systemic importance in the US financial system. Our approach can be used to monitor companies' systemic importance, enabling transparent macroprudential supervision.

JEL Classification: G01, G18, G32, G38, C21, C51, C63

1. Introduction

The financial crisis of 2007–09 has shown that cross-sectional dependencies between assets and credit exposure can cause risks of individual banks to cascade, ultimately substantially threatening the stability of the entire financial system.¹ Under certain economic conditions, company-specific risk cannot be appropriately assessed in isolation without accounting for potential risk spillover effects from other firms. Indeed, it is not merely the size

* We thank Thierry Foucault und two anonymous referees, as well as Tobias Adrian, Frank Betz, Christian Brownlees, Markus Brunnermeier, Jon Danielsson, Robert Engle, Tony Hall, Simone Manganelli, Robin Lumsdaine, Tuomas Peltonen as well as participants of the annual meeting of the Society for Financial Econometrics (SoFiE) in Chicago, the 2011 annual European Meeting of the Econometric Society in Oslo, the 2011 Humboldt-Copenhagen Conference in Copenhagen, the 2012 Financial Econometrics Conference in Toulouse and the European Central Bank high-level conference on “Financial Stability: Methodological Advances and Policy Issues”, Frankfurt, June 2012. Research support by the Deutsche Forschungsgemeinschaft via the Collaborative Research Center 649 “Economic Risk” is gratefully acknowledged. N. H. acknowledges research support by the Wiener Wissenschafts-, Forschungs- und Technologiefonds (WWTF). J. S. thanks the European Union Seventh Framework Programme (FP7-SSH/2007-2013, grant agreement 320270—SYRTO) for financial support.

¹ For a thorough description of the financial crisis, see, for example, Brunnermeier (2009).

and idiosyncratic risk of a firm but also its interconnectedness with other firms that determines its systemic relevance. The latter is a firm's potential to increase the risk of failure of the entire system—which we denote as systemic risk.² There is a broad consensus that any prudential regulatory policy should account for the consequences of network interdependencies in the financial system. In practice, however, any attempt at transparent implementation of such a policy must fail, as long as suitable empirical measures of firms' individual risk, risk spillovers, and systemic relevance are not available. In particular, how to quantify individual risk exposure and systemic risk contributions in an appropriate but parsimonious and empirically tractable way, given the prevailing underlying network structure, remains an open question. Moreover, there is need for empirically feasible measures that rely only on available data of publicly disclosed balance sheets and market information but still account for the complexity of the financial system.

A general empirical assessment of systemic relevance cannot build on the vast theoretical literature on financial network models and financial contagion, as such studies typically require detailed information on intra-bank asset and liability exposures (see, e.g., Allen and Gale, 2000; Freixas, Parigi, and Rochet, 2000; Leitner, 2005). Such data are generally not publicly disclosed, and even supervisory authorities can only collect partial information on inter-bank linkages. Available empirical studies linked to this literature can therefore only partially contribute to a full picture of companies' systemic relevance, as these studies focus on particular parts of specific markets at particular times under particular financial conditions (see, e.g., Furfine, 2003; Uppner and Worms, 2004, for Germany and the USA, respectively).³ Furthermore, assessing risk interconnections on the basis of multivariate failure probability distributions has proven to be statistically complex in the absence of restrictive assumptions (see, e.g., Boss *et al.*, 2004, or Zhou, 2009, and references therein). Finally, for banking regulators, it is often unclear how complex structures eventually translate into dynamic and predictable measures of systemic relevance.

The objective of this article is to develop an applicable measure of a firm's systemic relevance, explicitly accounting for the company's

² Bernanke (2009) and Rajan (2009) stress the danger induced by institutions that are “too interconnected to fail” or “too systemic to fail”, in contrast with firms that are simply “too big to fail”.

³ See also Cocco, Gomes, and Martins (2009) on parts of the financial sector in Portugal, Elsinger, Lehar, and Summer (2006) for Austria, and Degryse and Nguyen (2007) for Belgium. A rare exception is the unique data set for India with full information on the intra-banking market studied in Iyer and Peydríó (2011).

interconnectedness within the financial sector. We assess companies' risk of financial distress on the basis of share price information, publicly accessible market data, and balance sheet data. Our measure quantifies the risk of distress of individual companies and of the entire system, using tails of the corresponding asset return distributions. Consequently, it builds on extreme conditional quantiles and thus the concept of conditional Value-at-Risk (VaR), a popular and widely accepted measure of tail risk.⁴ For each firm, we identify its so-called relevant (tail) risk drivers as the set of macroeconomic fundamentals, firm-specific characteristics, and risk spillovers from other institutions driving the company's VaR. Such a conditional VaR specification yields a reliable measure of a firm's idiosyncratic risk in the presence of network effects. Moreover, detecting which firms an institution is connected with and measuring the strengths of these connections enables us to construct a tail risk network for the financial system. A company's contribution to systemic risk is then defined as the effect of an increase in its individual tail risk on the VaR of the financial system.

The underlying statistical setting is a two-stage quantile regression approach: in the first step, firm-specific VaRs are estimated as functions of firm characteristics, macroeconomic state variables, and tail risk spillovers of other banks captured as loss exceedances. The major challenge is to shrink the high-dimensional set of possible cross-linkages between all firms to a feasible number of relevant risk connections. We address this issue statistically as a model selection problem in individual institution's VaR specifications, a problem that we solve in the first step. Specifically, we use a novel Least Absolute Shrinkage and Selection Operator (LASSO) technique (see Belloni and Chernozhukov, 2011), which allows for identification of the relevant tail risk drivers for each company in a data-driven way. The resulting risk interconnections are represented in a network graph, as illustrated, for example, in Figure 1 for the system of the fifty-seven largest US financial companies. In the second step, to measure a firm's systemic impact, we individually estimate the VaR of a value-weighted index of the financial sector as a function of the firm's estimated VaR while controlling for the pre-identified company-specific risk drivers and the macroeconomic state variables. We derive standard errors that explicitly account for estimation errors that arise in the first estimation step. Additionally, we utilize bootstrap methods needed for accurate parameter tests in finite samples.

⁴ In principle, our methodology could also be adapted to other tail risk measures such as, for example, expected shortfall. Such a setting, however, would involve additional estimation steps and complications, probably inducing an overall loss of accuracy in our results, given the limited amount of available data.

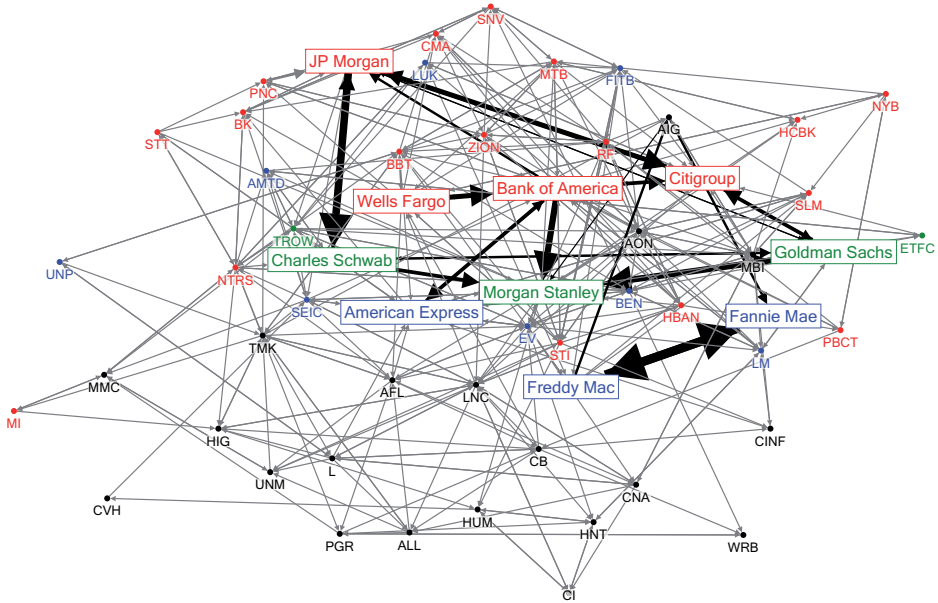


Figure 1. Risk network of the US financial system schematically highlighting key companies in the system in 2000–08. Details on all firms in the system that appear as unlabeled shaded nodes will be provided later in the article. An arrow pointing from firm j to firm i reflects the impact of extreme returns of j on the VaR of i (VaR^i), a connection that is identified as relevant through the statistical selection techniques presented in the remainder of the article. VaRs are measured in terms of 5% quantiles of the return distribution. The effect of j on i is measured in terms of the impact of an increase of the return X^j on VaR^i , given that X^i is below its 10% quantile, that is, i 's so-called loss exceedance. The size of the respective increase in VaR^i , given a 1% increase in the loss exceedance of i , is reflected in the thickness of the respective arrowhead, whereby we distinguish between three categories: thin arrowheads indicate an increase of up to 0.4, medium-sized arrowheads indicate an increase of 0.4–0.8, and thick arrowheads indicate an increase greater than 0.8. The thickness of the line of an arrow reflects these same categories. If an arrow points in both directions, the thickness of the line corresponds to the larger of the two effects. The graph is constructed so that the total length of all arrows in the system is minimized. Accordingly, more highly interconnected firms are located in the center.

We determine a company as systemically relevant if the marginal effect of the firm's VaR on the VaR of the system is statistically significant and nonnegative. In analogy to an (inverted) asset pricing relationship in quantiles, we call this marginal effect systemic risk beta. It is modeled as a function of firm-specific characteristics, such as leverage, maturity mismatch, and size, while controlling for macroeconomic conditions and the firm's network position. Thus, a firm's marginal systemic impact can change due to varying market or balance sheet conditions, although its

individual risk level might be identical at different points in time. The total increase in the system VaR due to a change in a firm-specific VaR is obtained as the product of the firm's systemic risk beta and its VaR. The latter, called the "realized" systemic risk beta, rises with increases in the firm's VaR. We use it to compare the levels of systemic importance of different companies and thus rank them across the system.

Our empirical results reveal a high degree of tail risk interconnectedness among US financial institutions, network effects that are dominant drivers of firms' individual risk. Detected spillover channels can be largely attributed to direct credit or liquidity exposure, although in some cases, they may also result from common factors, for example, factors specific to the sector or the business model, which are not covered by our firm-specific control variables. Generally, these links contain fundamental information for supervisory authorities but also for company risk managers. Based on the topology of the systemic risk network, we categorize firms into three broad groups, according to their type and extent of connectedness with other companies: main risk transmitters, risk recipients, and companies that both receive and transmit tail risk. From a supervisory point of view, the second group has the least systemic impact. Monitoring this group, however, may nevertheless convey important information on hidden risks and possible threats induced by a high degree of interconnectedness. The highest attention of supervisory authorities should be directed toward firms that mainly act as risk drivers or are highly interconnected risk transmitters within the system. These are firms, labeled "too interconnected to fail", in the center of the network, and also risk producers at the network periphery that are linked to only a few heavily connected risk transmitters.

While the systemic risk network yields qualitative information regarding risk channels and the roles of companies within the financial system, estimates of systemic risk betas allow us to quantify the systemic relevance of individual firms and thus complement the full picture. Ranking companies based on (realized) systemic risk betas show that large depositories are particularly risky. After controlling for relevant network effects, these firms have overall the strongest impact on systemic risk and should be regulated accordingly. Time series patterns of (realized) systemic risk betas indicate that most companies' systemic risk contributions sharply increased during the 2007–08 financial crisis, effects that were particularly pronounced for firms that experienced financial distress during the crisis and are (*ex post*) identified as clearly systemically risky under our approach. Figure 2 illustrates the paths of their marginal systemic contributions, as reflected in their systemic risk betas and their exposures to idiosyncratic tail risk, as quantified by their VaRs. A precrisis case study confirms the validity of our




Figure 2. Systemic importance of five exemplary firms in the US financial system at two points in time before and at the height of the financial crisis, 2008. Systemic relevance is determined by the statistical significance and positivity of “systemic risk betas” quantifying the marginal increase of the VaR of the system, given an increase in a bank’s VaR, while controlling for the bank’s (pre-identified) risk drivers. All VaRs are computed at the 5% level and are by definition positive. We depict the degree of systemic relevance by the size of the respective “realized” version of the systemic risk beta, that is, the product of the risk beta and the corresponding VaR of a company, representing the company’s total effect on systemic risk. Connecting lines are added to graphically highlight changes between the two points in time but do not represent actual evolutionary paths. The size of each element in the graph reflects the size of the VaR of the respective company at each of the two points in time. We use the following scale: the element is k times the standard size with $k=1$ for $\text{VaR} \leq 0.05$, $k=1.5$ for $\text{VaR} \in (0.05, 0.1]$, $k=2$ for $\text{VaR} \in (0.1, 0.15]$, $k=3$ for $\text{VaR} \in (0.2, 0.25]$ and $k=5.5$ for $\text{VaR} \in (0.65, 0.7]$. Attached numbers inside the figure mark the position of the respective company in an overall ranking of the fifty-seven largest US financial companies for each of the two time points.

methodology, as firms such as, for example, Lehman Brothers are *ex-ante* identified as highly systemically relevant. It is well-known that the subsequent failure of this firm indeed had a huge impact on the stability of the entire financial system. Similarly, the extensive bailouts of American International Group (AIG), Freddie Mac, and Fannie Mae can be justified, given their high systemic risk betas and high interconnectedness as of the end of 2007.

The remainder of the article is structured as follows. Section 2 describes the article’s links with related literature and presents the underlying data. In Section 3, we present the model and the procedure used to estimate individual companies’ VaRs, which are the basis for determining the systemic tail risk network structure. The notion of a realized systemic risk beta is formally introduced in Section 4, and realized systemic risk betas are identified for

each firm in an individually tailored parsimonious partial equilibrium setting. This section also presents the corresponding estimation procedure and valid inference for a two-step quantile regression setting. Our empirical results are presented in the form of systemic risk rankings. In Section 5, we validate our model and results. In particular, in a case study that uses only precrisis data, we illustrate that realized systemic risk betas are effective in predicting the distress and systemic relevance of five large financial institutions that were affected by the financial crisis. Section 6 concludes the article.

2. Literature and Data

2.1 RELATION TO THE RECENT EMPIRICAL LITERATURE

Our article relates to several strands of recent empirical literature on systemic risk contributions. Building on the concept of VaR, Adrian and Brunnermeier (2011) were the first to model systemic risk contributions based on balance sheet characteristics. They introduce the so-called CoVaR as a firm's (*conditional*) VaR, given that some other firm's stock return takes a certain benchmark value (e.g., the individual VaR). There are, however, substantial conceptual differences between their approach and ours: the realized systemic risk beta in our approach is the direct marginal effect of an individual VaR on the VaR of the system. Conversely, CoVaR builds on the marginal effect of the return and is only evaluated at the value of the (pre-estimated) VaR. As returns are below their VaR(q) ($1 - q$)% of the time,⁵ the estimated marginal systemic importance of CoVaR tends to systematically overrate firms with lower average returns for identical risk levels. Furthermore, CoVaR can by definition only vary over time through the channel of individual VaRs. Due to multicollinearity, however, it cannot be modeled in terms of firm-specific variables. Thus, changes in firms' systemic relevance only result from variations in underlying macroeconomic factors, while variations in firms' leverage and interdependence with other institutions have no direct effect. Under our approach, in contrast, we identify network spillovers as crucial elements in measuring individual risk and in unbiased estimation of systemic relevance. This is illustrated in a robustness study in Subsection 3.3.a. Moreover, the proposed realized

⁵ The VaR(q) is defined as the negative q -conditional return quantile.

systemic risk beta captures variations in firms' marginal systemic importance driven by changes in firm-specific characteristics.

Our work also complements papers, such as Acharya *et al.* (2010), Brownlees and Engle (2012) and Acharya, Engle, and Richardson (2012), which measure a company's systemic relevance in terms of the size of potential bailout costs. Such approaches cannot detect spillover effects driven by the topology of the risk network and may tend to underestimate the systemic importance of highly interconnected companies. While Brownlees and Engle (2012) study an individual firm's conditionally expected asset return given distress of the system, we investigate the reverse relation and measure the effect on the system given a firm is in financial trouble. Taking complementary perspectives, the two approaches measure different dimensions of systemic risk. However, as our model is based on economic state variables and loss exceedances, it automatically adjusts and prevails in distress scenarios under external shocks. This is a clear advantage of our approach compared with pure time series approaches (cp. e.g., White, Kim, and Manganello, 2010; Brownlees and Engle, 2012). As illustrated in the validity case study in Section 5.2, the estimated systemic risk betas indicate an increase in systemic relevance of some companies earlier than in competing settings.

Our work also augments research of Billio *et al.* (2012), who present a collection of different systemic risk measures. These measures mainly build on regressions of (conditional) means of returns. However, assessing and predicting systemic and firm-specific risk requires regression in the (left) tails of asset return distributions rather than the center. Hence, our approach focuses on extreme quantiles and thus substantially differs from a correlation type analysis, as in Billio *et al.* (2012). Moreover, in contrast to our approach, the latter authors' determination of causality is based only on pairwise relations. Such a setting, however, produces misleading results in a high-dimensional interconnected system, as it is impossible to identify whether one firm drives another or if both are driven by a third company. Our results are also complementary to network analysis based on volatility spillovers in vector autoregressive systems, for example, Diebold and Yilmaz (2012) and Diebold and Yilmaz (2013).

Finally, we contribute to macroeconomic approaches that take a more aggregated view, for example, the literature on systemic risk indicators (e.g., Segoviano and Goodhart, 2009; Giesecke and Kim, 2011) or papers on early warning signals (e.g., Koopman, Lucas, and Schwaab, 2011; Schwaab, Koopman, and Lucas, 2011).

2.2 DATA

Our analysis focuses on publicly traded US financial institutions. The list of included companies in Table I comprises depositories, broker dealers, insurance companies, and other firms.⁶ To assess a firm's systemic relevance, we use publicly accessible market and balance sheet data. The forward-looking nature and real-time availability of equity market data serves well in providing an immediate and transparent measure of systemic risk. The advantage of timeliness is evident, even if new financial regulations compel institutions to reveal information on mutual credit linkages and leverage to supervisory authorities. Currently, however, data on connections between firms' assets and obligations are largely proprietary and far from comprehensive, even for supervisors.

Daily equity prices are obtained from Datastream and are converted to weekly log returns. To account for the general state of the economy, we use weekly observations of seven lagged macroeconomic variables, M_{t-1} , as suggested and used by Adrian and Brunnermeier (2011) (abbreviations used in the remainder of the article are given in brackets): the implied volatility index, VIX, as computed by the Chicago Board Options Exchange (vix), a short-term "liquidity spread", computed as the difference of the 3-month collateral repo rate (available on Bloomberg) and the 3-month Treasury bill rate from the Federal Reserve Bank of New York (repo), the change in the 3-month Treasury bill rate (yield3m) and the change in the slope of the yield curve, corresponding to the spread between the 10-year and 3-month Treasury bill rate (term). Additionally, we utilize changes in credit spreads between BAA rated bonds and the Treasury bill rate (both at 10-year maturity) (credit), the weekly equity market return from CRSP (marketret) and the 1-year cumulative real estate sector return, computed as the value-weighted average of real estate companies, available in the CRSP data base (housing).⁷ Analyzing the time series properties of the variables reveals that, with the exceptions of vix and housing, they are stationary. Applying the Engle and Granger (1987) two-step procedure, however, we find evidence of cointegration between the two variables, which implies that their joint explanation in the model is stationary and that inference thus remains valid (see Pagan and Wickens, 1989). Therefore, to maintain the

⁶ Companies are classified into these groups according to their two-digit SIC codes, following the categorization in Adrian and Brunnermeier (2011), Appendix C.

⁷ We found that this set of aggregate financial market variables provides sufficient explanatory power that is not further increased by additional controls such as, for example, Fama-French type factors (see Subsection 3.3.a. for details).

Table 1. Included financial institutions in alphabetical order within sectors

^aIncluded only in the case study.

Depositories (21)	Others (11)	Insurance Comp. (20)
BB T Corp. (BBT)	American Express Co (AXP)	AFLAC Inc. (AFL)
Bank of New York Mellon (BK)	Eaton Vance Corp. (EV)	Allstate Corp. (ALL)
Bank of America Corp. (BAC)	Fed. Home Loan Mortg. Corp. (FRE)	American International Group (AIG)
Citigroup Inc. (C)	Fed. National Mortgage Assn (FNM)	AON Corp. (AON)
Comerica Inc. (CMA)	Fifth Third Bancorp. (FITB)	Berkley WR Corp. (WRB)
Hudson City Bancorp. Inc. (HCBK)	Franklin Resources Inc. (BEN)	CIGNA Corp. (CI)
Huntington Bancshares Inc. (HBAN)	Legg Mason Inc. (LM)	C N A Financial Corp. (CNA)
JP Morgan Chase & Co (JPM)	Leucadia National Corp. (LUK)	Chubb Corp. (CB)
M & T Bank Corp. (MTB)	SEI Investments Company (SEIC)	Cincinnati Financial Corp. (CINF)
Marshall & Ilsley Corp. (MI)	TD Ameritrade Holding Corp. (AMTD)	Coventry Health Care Inc. (CVH)
NY Community Bancorp. (NYB)	Union Pacific Corp. (UNP)	Hartford Financial (HIG)
Northern Trust Corp. (NTRS)		HEALTH NET INC (HNT)
Peoples United Financial Inc. (PBCT)	Broker-Dealers (7)	Humana Inc. (HUM)
PNC Financial Services Group (PNC)	E Trade Financial Corp. (ETFC)	Lincoln National Corp. (LNC)
Financial Corp. New (RF)	Goldman Sachs Group Inc. (GS)	Loews Corp. (L)
S L M Corp.	Lehman Brothers (LEH) ^a	Marsh & McLennan Inc. (MMC)
State Street Corp. (STT)	Merrill Lynch (ML) ^a	MBIA Inc. (MBI)
Suntrust Banks Inc. (STI)	Morgan Stanley Dean Witter & Co (MS)	Progressive Corp. Ohio (PGR)
Synovus Financial Corp. (SNV)	Schwab Charles Corp. New (SCHW)	Torchmark Corp. (TMK)
Wells Fargo & Co (WFC)	T Rowe Price Group Inc. (TROW)	Unum Group (UNM)
Zions Bancorp. (ZION)		

comparability of our results with those in the literature, we use the two regressors in levels.

To capture characteristics of individual institutions that predict a bank’s propensity to become financially distressed, C_{t-1}^i , we follow Adrian and Brunnermeier (2011) and use (i) leverage, calculated as the value of total assets divided by total equity (in book values) (LEV); (ii) maturity mismatch, measuring short-term refinancing risk, calculated as short-term debt net of cash divided by total liabilities (MMM); (iii) market-to-book value, defined as the ratio of the market value to the book value of total equity (BM); (iv) market capitalization, defined as the logarithm of market valued total assets (SIZE); and (v) equity return volatility, computed from daily equity return data (VOL). The system return is chosen as the return on the financial sector index provided by Datastream. It is computed as the value-weighted average of prices of 190 US financial institutions.⁸

As balance sheets are available only on a quarterly basis, we interpolate the quarterly data to daily level, using cubic splines, and then aggregate them back to calendar weeks.⁹ We focus on fifty-seven financial institutions that existed throughout the period from the beginning of 2000 to the end of 2008, resulting in 467 weekly observations of individual returns. This criterion excludes companies that defaulted during the financial crisis. The latter are analyzed separately in a shorter sample case study.

3. A Tail Risk Network

3.1 DETERMINING DRIVERS OF FIRM-SPECIFIC TAIL RISK

We measure the tail risk of a company with asset return X_t^i at time t as its conditional Value-at-Risk (VaR), $\text{VaR}_{q,t}^i$, given a set of company-specific tail risk drivers $\mathbf{W}_t^{(i)}$:

$$\Pr(-X_t^i \geq \text{VaR}_{q,t}^i | \mathbf{W}_t^{(i)}) = \Pr(X_t^i \leq Q_{q,t}^i | \mathbf{W}_t^{(i)}) = q \tag{1}$$

with $\text{VaR}_{q,t}^i = \text{VaR}_{q,t}^i(\mathbf{W}_t^{(i)}) = -Q_{q,t}^i$ denoting the (negative) conditional q -quantile of X_t^i .¹⁰ The relevant i -specific tail risk drivers are determined

⁸ See Adrian and Brunnermeier (2011), Appendix C, who explicitly show that this variable induces no inherent endogeneity in the model.

⁹ For in-sample estimation, this interpolation step captures changes in balance sheet characteristics in a smoother way than the use of plain data. For forecasting purposes, however, interpolation is not possible. See Hautsch, Schaumburg, and Schienle (2014) for details.

¹⁰ Defining VaR as the *negative* p -quantile ensures that the VaR is positive and is interpreted as a loss position.

out of a large set of potential regressors \mathbf{W}_t containing lagged macroeconomic state variables \mathbf{M}_{t-1} , lagged firm-specific characteristics \mathbf{C}_{t-1}^i , the i -specific lagged return X_{t-1}^i and influences of companies other than i , $\mathbf{E}_t^{-i} = (E_t^j)_{j \neq i}$. We capture these network dependencies in terms of so-called loss exceedances, defined (for firm j) as $E_t^j = X_t^j \mathbf{1}(X_t^j \leq \hat{Q}_{0.1}^j)$, where $\hat{Q}_{0.1}^j$ is the unconditional 10% sample quantile of X^j . Hence, company j only affects the VaR of company i if the former is in distress.

We model the conditional VaR of firm i at time $t = 1, \dots, T$ as a linear function of the i -specific tail risk drivers $\mathbf{W}_t^{(i)}$,

$$\text{VaR}_q^i = \mathbf{W}_t^{(i)'} \xi_q^i. \tag{2}$$

This relation could be estimated using a corresponding linear model in the corresponding return quantile

$$X_t^i = -\mathbf{W}_t^{(i)'} \xi_q^i + \varepsilon_t^i, \quad \text{with} \quad Q_q(\varepsilon_t^i | \mathbf{W}_t^{(i)}) = 0 \tag{3}$$

if we knew the i -relevant risk drivers $\mathbf{W}^{(i)}$ selected from \mathbf{W} . Then the estimates $\widehat{\xi}_q^i$ of ξ_q^i could be obtained from the standard linear quantile regression (Koenker and Bassett, 1978) by minimizing

$$\frac{1}{T} \sum_{t=1}^T \rho_q \left(X_t^i + \mathbf{W}_t^{(i)'} \xi_q^i \right) \tag{4}$$

with loss function $\rho_q(u) = u(q - I(u < 0))$, where the indicator $I(\cdot)$ is 1 for $u < 0$ and zero otherwise, and

$$\widehat{\text{VaR}}_{q,t}^i = \mathbf{W}_t^{(i)'} \widehat{\xi}_q^i. \tag{5}$$

The relevant risk drivers $\mathbf{W}^{(i)}$ for firm i , however, are unknown and must be determined from \mathbf{W} in advance. Appropriate model selection techniques are not straightforward in the given setting, as tests of the individual significance of single variables do not account for the (possibly high) collinearity between the covariates. Similarly, sequences of joint significance tests have too many possible variations to be easily checked in cases of more than sixty variables. Therefore, we choose the relevant covariates in a data-driven way by employing a statistical shrinkage technique known as the least absolute shrinkage and selection operator (LASSO). LASSO methods are standard for high-dimensional conditional mean regression problems (see Tibshirani, 1996) and have recently been adapted to quantile regression by Belloni

and Chernozhukov (2011). Accordingly, we run an l_1 -penalized quantile regression and calculate, for a fixed individual penalty parameter λ^i ,

$$\tilde{\xi}_q^i = \operatorname{argmin}_{\xi^i} \frac{1}{T} \sum_{t=1}^T \rho_q(X_t^i + \mathbf{W}'_t \xi^i) + \lambda^i \frac{\sqrt{q(1-q)}}{T} \sum_{k=1}^K \hat{\sigma}_k |\xi_k^i|, \tag{6}$$

with the set of potentially relevant regressors $\mathbf{W}_t = (W_{t,k})_{k=1}^K$, which are demeaned, componentwise variation $\hat{\sigma}_k^2 = \frac{1}{T} \sum_{t=1}^T (W_{t,k})^2$, and the loss function ρ_q as in Equation (4). The key idea is to select relevant regressors according to the absolute values of their estimated marginal effects (scaled by the regressor’s variation) in the penalized VaR regression (6). Regressors are eliminated if their shrunken coefficients are sufficiently close to zero. Here, all firms in \mathbf{W} with absolute marginal effects $|\tilde{\xi}_q^i|$ below a threshold $\tau = 0.0001$ are excluded, and only the $K(i)$ remaining relevant regressors $\mathbf{W}^{(i)}$ are retained. Hence, LASSO deselects regressors that contribute only small amounts of variation. Due to the additional penalty term in (6), all coefficients $\tilde{\xi}_q^i$ are generally downwardly biased in finite samples. Therefore, we reestimate the unrestricted model (4) only with the selected relevant regressors $\mathbf{W}^{(i)}$, yielding the final estimates $\hat{\xi}_q^i$. This post-LASSO step produces finite sample estimates of the coefficients ξ_q^i , estimates that are superior to the original LASSO estimates or plain quantile regression results without penalization, which suffer from overidentification problems (see the original paper by Belloni and Chernozhukov, 2011 for the consistency proof of the post-LASSO step).

The selection of relevant risk drivers via LASSO crucially depends on the choice of the company-specific penalty parameter λ^i . The larger is the chosen value of λ^i , the more regressors are eliminated. Conversely, in case of $\lambda^i = 0$, we are back in the standard quantile regression setting (4) without any deselection. For each institution, we determine the appropriate penalty level λ^i in a completely data-driven way by using the supremum norm of a rescaled gradient of the sample criterion function, evaluated at the true parameter value, as in Belloni and Chernozhukov (2011).¹¹ Consequently, the number and set of relevant risk drivers are determined only from the data, without any restrictive pre-assumptions. For further details on this empirical procedure, see (A.5) in the Appendix.

¹¹ See Step 1 for (A.5) in the Appendix for the scaling and the exact formula.

Evaluating the goodness of fit of the resulting conditional VaR specifications requires quantifying how well the model captures the specific percentile of the return distribution and how well the model predicts the size and frequency of losses. With respect to the latter issue, it is not sufficient to use a simple quantile-based modification of the conventional R^2 statistic. We therefore consider a VaR specification as inadequate if it either fails to produce the correct empirical level of VaR exceedances or the sequence of exceedances is not independently and identically distributed. This ensures that VaR violations today do not contain information about VaR violations in the future and that both occur according to the same distribution. This can be formally tested using a likelihood ratio (LR) version of the dynamic quantile (DQ) test developed in Engle and Manganelli (2004) which is described in detail in (A.7) in the Appendix. Berkowitz, Christoffersen, and Pelletier (2011) show that this LR test has superior size and power properties compared with competing conditional VaR backtesting methods, which dominate plain unconditional level tests (as e.g., Kupiec (1995)).

Using the LASSO selection procedure described above, we estimate VaR specifications for $q = 0.05$ for all individual companies.¹² Table II provides exemplary VaR^i (post-)LASSO regression results for firms in four industry sectors: depositories, insurance companies, broker dealers, and others. We find that the dominant drivers of company-specific VaRs are loss exceedances of other firms. In their presence, macroeconomic variables and firm-specific characteristics often do not have any statistically significant influence and are not selected by the LASSO procedure. In Table II, for instance, VaR specifications for Goldman Sachs (GS), Morgan Stanley (MS), JP Morgan (JPM), and AIG exclusively contain loss exceedances of other firms. The importance of cross-firm effects as drivers of individual tail risk is confirmed by a joint significance test of the individually selected loss exceedances E_t^{-i} and the superiority of resulting VaR forecasts. The latter aspect is analyzed in Subsection 3.3.

As a result of our estimation procedure, we not only detect “relevant” risk connections but can assign directions thereof. Selecting E^j as a relevant risk driver of VaR^i implies a directed link from j to i . If, in addition, E^i significantly affects VaR^j , we observe a bidirectional relation, which is, however, not symmetric.¹³ Note that our analysis, thus, is not affected by simultaneity

¹² Due to the limited number of observations, we refrain from considering more extreme probabilities.

¹³ For the significance of effects, see Subsection 3.3.a.

Table II. Exemplary post-LASSO quantile regressions for VaR^j with $q = 0.05$

Regressors were selected by LASSO as outlined in Section 3.1. Ex. j is the loss exceedance of company j , all other regressors are as in Section 2.2.

	Value	Std. Error	t -ratio	p -value
GS				
(Intercept)	-0.046	0.004	-12.54	0.000
Ex.C	-0.239	0.205	-1.17	0.243
Ex.JPM	-0.014	0.119	-0.121	0.904
Ex.LM	-0.215	0.111	-1.932	0.054
Ex.MS	-0.403	0.079	-5.096	0.000
Ex.SCHW	-0.282	0.244	-1.153	0.249
MS				
(Intercept)	-0.041	0.003	-16.017	0.000
Ex.AIG	-0.106	0.026	-4.036	0.000
Ex.AON	0.445	0.145	3.066	0.002
Ex.BAC	-0.604	0.145	-4.157	0.000
Ex.EV	-0.158	0.134	-1.179	0.239
Ex.GS	-0.634	0.121	-5.236	0.000
Ex.HBAN	-0.273	0.136	-2.006	0.045
Ex.HCBK	-0.452	0.28	-1.611	0.108
Ex.MTB	-0.269	0.193	-1.392	0.165
Ex.SCHW	-0.381	0.116	-3.294	0.001
Ex.SEIC	-0.229	0.154	-1.485	0.138
Ex.STT	-0.174	0.176	-0.986	0.325
Regions Financial				
(Intercept)	-0.004	0.004	-1.072	0.284
Ex.AMTD	-0.091	0.04	-2.274	0.023
Ex.AON	-0.256	0.086	-2.998	0.003
Ex.BBT	-0.307	0.104	-2.95	0.003
Ex.FITB	0.032	0.087	0.37	0.712
Ex.HBAN	-0.042	0.064	-0.661	0.509
Ex.PBCT	-0.307	0.085	-3.598	0
Ex.STI	-0.244	0.114	-2.137	0.033
Ex.ZION	-0.196	0.1	-1.947	0.052
BM	0.024	0.007	3.221	0.001
VOL	0.251	0.16	1.568	0.118
Fannie Mae				
(Intercept)	-0.049	0.003	-17.075	0.000
Ex.AIG	-0.227	0.231	-0.981	0.327
Ex.FRE	-1.007	0.121	-8.298	0.000
AIG				
(Intercept)	-0.043	0.003	-14.026	0.000
Ex.FRE	-0.201	0.014	-14.033	0.000
Ex.MBI	-0.336	0.138	-2.423	0.016
Ex.RF	-0.455	0.051	-8.975	0.000
Ex.TMK	-0.813	0.721	-1.127	0.260
Torchmark				
(Intercept)	-0.019	0.003	-7.203	0

(continued)

Table II. (Continued)

	Value	Std. Error	<i>t</i> -ratio	<i>p</i> -value
Ex.AFL	-0.332	0.169	-1.962	0.05
Ex.ALL	-0.256	0.207	-1.237	0.217
Ex.BBT	-0.296	0.223	-1.329	0.184
Ex.HIG	-0.084	0.175	-0.483	0.63
Ex.LNC	0.002	0.135	0.018	0.986
Ex.NTRS	-0.002	0.115	-0.015	0.988
Ex.SEIC	-0.243	0.12	-2.023	0.044
Ex.UNM	-0.088	0.179	-0.489	0.625
Ex.UNP	-0.242	0.242	-1	0.318
repo	0.031	0.017	1.78	0.076
JP Morgan				
(Intercept)	-0.040	0.003	-12.963	0.000
Ex.BAC	-0.229	0.133	-1.724	0.085
Ex.BK	-0.237	0.129	-1.842	0.066
Ex.C	-0.380	0.22	-1.729	0.084
Ex.GS	-0.253	0.154	-1.648	0.199
Ex.PNC	-0.274	0.077	-3.583	0.000
Ex.SCHW	-0.410	0.118	-3.472	0.001
American Express				
(Intercept)	-0.035	0.003	-11.723	0
Ex.AFL	-0.42	0.408	-1.03	0.303
Ex.BAC	-0.361	0.205	-1.757	0.08
Ex.BBT	-0.145	0.126	-1.151	0.25
Ex.BEN	-0.112	0.139	-0.808	0.42
Ex.CINF	-0.153	0.153	-0.999	0.318
Ex.EV	-0.181	0.163	-1.112	0.267
Ex.L	0.014	0.114	0.122	0.903
Ex.SEIC	-0.106	0.09	-1.186	0.236
Ex.SLM	0.073	0.067	1.09	0.276
Ex.STT	-0.351	0.159	-2.2	0.028
Ex.TROW	-0.3	0.126	-2.39	0.017

biases. For instance, a highly negative return of company j increases the conditional loss quantile and thus the VaR of firm i . The latter, however, does not necessarily imply a higher (realized) loss exceedance of i , as the relationship between a specific conditional quantile and the conditional distribution of exceedances (for a fixed threshold) is generally unknown. Even if quantiles and exceedances are positively related, an increase in VaR^i may only induce an increase in the expected loss exceedance, not necessarily in the realized loss exceedance. Consequently, the potential effect of a simultaneity bias (if it exists at all) is expected to be much weaker than in classical mean regressions and thus can be safely ignored. We consider it an advantage of our approach that it addresses network dependencies in a parsimonious way,

avoiding infeasible treatment of an explicit large system of conditional quantiles.¹⁴

There may be several economic reasons for linkages between two companies—reasons, however, that cannot be empirically identified from publicly disclosed market data.¹⁵ By including firm-specific characteristics and macroeconomic state variables in our model, we do, however, prevent that determined identified risk connections result from common (risk) factors. Hence, we rule out the possibility that tail dependencies are driven, for instance, by periods of high volatility, the flattening of yield curves or declining overall credit quality. Accordingly, the identified risk connections are likely to be attributable to remaining factors, such as credit or liquidity exposure, business model commonalities, or sector-specific risk factors. In this sense, connections between close competitors such as GS and MS or the influence of the mortgage company Freddie Mac (FRE) on AIG confirm market evidence.

3.2 NETWORK MODEL AND STRUCTURE

We construct a tail risk network of the system, using individually selected loss exceedances.¹⁶ Taking all firms as nodes in such a network, there is an influence of firm j on firm i , if E^j is LASSO-selected in (6) as a relevant driver of VaR_q^i . Let E^j be the k -th component of $\mathbf{W}^{(i)}$. Then the corresponding coefficient $\xi_{q,k}^i$ in ξ_q^i marks the impact of firm j on firm i in the network. If E^j is not selected as a relevant risk driver of firm i , there is no network arrow from firm j to firm i .

An overview of the identified tail risk connections between all companies (based on VaR specifications with $q = 0.05$) is provided in Table III.¹⁷

¹⁴ Statistically, it is an open question how to generally handle such a system of conditional quantiles. In contrast to relations in (conditional) means, it is unclear how marginal q -quantiles constitute the corresponding quantile in the joint distribution under appropriate independence assumptions. Only in lags, restricted to very small dimensions and under strong assumptions, have solutions been obtained via CaViAR type structures (see White, Kim, and Manganelli, 2010).

¹⁵ Note that a valid empirical classification into different types of linkages would require comprehensive data on the credit and liquidity exposures of firms. Such information, however, is largely not publicly available.

¹⁶ In the Bayesian network literature, a network that builds on direct one-step influences constitutes a so-called Markov blanket, which is assumed to contain all relevant information needed to predict a node's role in the network (see Friedman, Geiger, and Goldszmidt, 1997).

¹⁷ More extreme probabilities are theoretically feasible but require a larger number of observations for sufficient statistical precision. We also used alternative thresholds for

Table III. Tail risk cross dependencies

For each company i , we list direct risk drivers and risk recipients within the network topology. The risk drivers are loss exceedances selected by the LASSO technique (6) as “relevant” regressors for the VaR^i model ($q = 0.05$) (“Influencing companies”). Direct risk recipients (“Influenced companies”) are companies for which the loss exceedance of company i appears as relevant via LASSO in their corresponding VaR^i .

Name	Influencing companies	Influenced companies
Broker dealers		
ETFC	AMTD,GS,MS	AMTD,C
GS	C,JPM,LM,MS,SCHW	BEN,C,ETFC,JPM,LM,MS,SCHW
MS	AIG,AON,BAC,EV,GS,HBAN,HCBK,MTB,SCHW,SEIC,STT	AMTD,BAC,EV,GS,HUM,LMC,ETFC,SEIC
SCHW	AMTD,GS,JPM,NTRS,TROW	AMTD,MS,GS,JPM
TROW	AMTD,BEN,EV,JPM,LUK,NTRS,SEIC,SNV	AON,MBI,MMC,AMTD,AXP,BEN,EV,NTRS,SCHW
Depositories		
BAC	AON,AXP,C,HBAN,LM,MS,MTB,PBCT,PNC,SEIC,STI,WFC	AXP,BBT,C,CMA,HCBK,JPM,LM,MBI,MS,MTB,PNC,STI,WFC
BBT	BAC,FITB,MTB,NTRS,STI,TMK,UNP,WFC	AXP,BEN,CMA,FRE,MTB,RF,TMK,UNP,WFC,ZION
BK	AXP,JPM,MTB,NTRS,SNV,STT,WFC	CMA,JPM,NTRS,SEIC,SNV
C	BAC,ETFC,FITB,GS,JPM,LNC,LUK,MBI,MTB	BAC,GS,JPM,LUK
CMA	AON,BAC,BBT,BK,HBAN,RF,SNV,WFC	AON,PNC,SNV,ZION
HBAN	AON,LNC,RF,STI,ZION	AON,BAC,CMA,EV,LNC,MS,PBCT,RF,ZION
HCBK	AON,BAC,MBI,MTB,NYB	MS,MTB
JPM	BAC,BK,C,GS,PNC,SCHW	BK,C,GS,SCHW,SEIC,TROW
MI	MMC,TMK	HIG,MMC
MTB	BAC,BBT,HCBK,NYB,SNV,ZION	AON,BAC,BBT,BK,HCBK,MS,SNV,WFC,ZION,C
NTRS	BEN,BK,LUK,MMC,SEIC,STT,TROW	AFL,AMTD,BBT,BEN,BK,HIG,MMC,PCR,SCHW,TMK,TROW,LUK,STT
NYB	PBCT,WFC	MTB,SLM,WFC,HCBK,PBCT
PBCT	HBAN,NYB	AON,BAC,CB,NYB,RF
PNC	BAC,CMA,STT,TMK,WFC,ZION	BAC,JPM,ZION
RF	AMTD,AON,BBT,FITB,HBAN,PBCT,STI,ZION	AIG,AON,CMA,EV,FITB,HBAN,MBI,SNV,STI,ZION
SLM	AON,AXP,FRE,MBI,NYB	AON,AXP,BEN,EV,FITB,MBI
SNV	BK,CMA,FITB,MTB,RF,ZION	BEN,BK,CMA,FITB,MTB,TROW
STI	AON,BAC,FITB,LNC,RF,WFC,ZION	AFL,AON,BAC,BBT,FITB,HBAN,RF,ZION,CINF,HUM,UNM,WFC
STT	AXP,NTRS	AXP,BK,NTRS,PNC,MS
WFC	BAC,BBT,CB,LNC,MTB,NYB,STI	FITB,PNC,STI,AFL,BAC,BBT,BK,CMA,NYB
ZION	BBT,CMA,HBAN,MTB,PNC,RF,STI	AON,RF,FITB,HBAN,LNC,MTB,PNC,SNV,STI

(continued)

Insurance companies	
AFL	ALL,AON,CNA,EV,NTRS,SEIC,STI,TMK,WFC
AIG	FRE,MBI,RF,TMK
ALL	CB,CNA,L,LNC,TMK
AON	CMA,HBAN,MBI,MTB,PBCT,RF,SLM,STI,TROW,ZION
CB	AFL,L,LNC,PBCT,PGR
CI	CNA,HNT,HUM,LNC
CINF	CB,MBI,STI
CNA	EV,L,LNC,MBI
CVH	HUM
HIG	CB,L,LNC,MI,NTRS,TMK
HNT	CI,EV,HUM,LM,LNC,PGR
HUM	CI,HIG,HNT,MS,STI
L	CB,CNA,LNC,TMK,UNP
LNC	CI,CNA,EV,HBAN,HIG,L,MS,SEIC,TMK,ZION
MBI	AIG,AON,BAC,BEN,CNA,FRE,RF,SLM,TROW
MMC	MI,NTRS,PGR,SEIC,TROW,UNM
PGR	AFL,ALL,NTRS,WRB
TMK	AFL,ALL,BBT,HIG,LNC,NTRS,SEIC,UNM,UNP
UNM	AFL,ALL,L,LNC,MMC,STI
WRB	BEN,CB,PGR
Others	
AMTD	ETFC,MS,NTRS,SCHW,SEIC,TROW
AXP	AFL,BAC,BBT,BEN,CINF,EV,L,SEIC,SLM,STT,TROW
BEN	AON,AXP,BBT,EV,GS,LM,MBI,NTRS,SLM,SNV,TROW
EV	AFL,AON,AXP,BEN,CB,HBAN,MS,RF,SEIC,SLM,TMK,TROW
FITB	AON,LUK,RF,SLM,SNV,STI,WFC,ZION
FNM	AIG,FRE
FRE	BBT,EV,FITB,FNM,LUK
LM	AON,BAC,BEN,CINF,EV,GS,HNT,MBI
LUK	C,LM,NTRS
SEIC	BK,CVH,JPM,TLN,C,MS
UNP	BBT,L
	AXP,CB,EV,PGR,TMK,UNM
	FNM,MBI,MS
	AFL,PGR,TMK,UNM
	AFL,BAC,BEN,CMA,EV,FITB,HBAN,HCBK,LM,MBI,MS,RF,SLM,STI
	ALL,CINF,EV,HIG,L,WFC,WRB
	HNT,HUM,LNC
	AXP,LM
	AFL,ALL,CI,L,LNC,MBI
	SEIC
	HUM,LNC,TMK
	CI,HUM,LM
	CI,HNT
	ALL,AXP,CB,CNA,HIG,LNC,UNM,UNP
	ALL,C,CB,CNA,HBAN,HIG,HNT,L,SEIC,STI,TMK,UNM,WFC,CI
	AIG,AON,BEN,C,CINF,HCBK,SLM,CNA,LM
	MI,NTRS,UNM
	MMC,CB,HNT,WRB
	AFL,BBT,EV,L,LNC,MI,PNC,AIG,ALL,HIG
	TMK,MMC
	PGR
	ETFC,RF,SCHW,TROW
	BAC,BEN,BK,EV,SLM,STT
	AXP,EV,LM,MBI,NTRS,TROW,WRB
	AFL,AXP,BEN,CNA,FRE,HNT,LM,LNC,MS,TROW
	BBT,C,FRE,RF,SNV,STI
	FRE
	AIG,MBI,SLM,FNM
	BAC,BEN,GS,HNT,LUK
	C,FITB,FRE,NTRS,TROW
	AFL,AMTD,AXP,BAC,EV,LNC,MMC,MS,NTRS,TMK,TROW
	BBT,TMK,L

The number of risk connections substantially varies over the cross-section of companies. To effectively illustrate identified risk connections and directions, we graphically depict the resulting network of companies in Figure 3. The layout of the network is chosen so that the sum of cross-firm distances is minimized. Consequently, the most highly connected firms are located in the center.

The resulting network topology allows us to distinguish between three major categories of firms: the first group contains companies with only a few incoming arrows but many outgoing ones; hence, such firms mainly act as risk drivers within the system. These are institutions whose failure may affect many others, while they themselves would be relatively unaffected by the distress of others. Such firms should be closely monitored by supervisory authorities, as the failure of such a bank could have widespread consequences. An example of such a bank is State Street Corporation (STT), one of the top ten US banks and a firm whose failure would affect financial services companies such as American Express and Northern Trust (NTRS), Bank of New York Mellon, and Morgan Stanley. Another example is the financial services firm SEI Investments, which has links to various large institutions such as Bank of America, American Express, Morgan Stanley, and the online broker TD Ameritrade (AMTD).

The second group contains companies that mainly act as risk takers. These companies are not necessarily systemically risky, but they may suffer severely from the distress of others and should account for such spillover effects in their internal risk management. According to Table III and Figure 3, such firms are primarily insurance companies.

The third group is the largest category. It consists of companies that serve as both risk recipients and risk transmitters that amplify tail risk spillovers by further disseminating risk into new channels. Due to their role as risk distributors, such companies are key systemic players and should be supervised accordingly. Examples of strongly connected companies in this category are GS, Citigroup, Morgan Stanley, AON Corporation (AON), Bank of America, American Express, and Freddie Mac. The latter was particularly affected by the 2008 credit crunch in the mortgage sector. Details on the specific roles of Citigroup and MS within the system are highlighted in Figure 4. Examples of firms with risk connections with only a few other institutions are Fannie Mae and AIG. Fannie Mae exhibits significant bilateral risk connections with its main competitor Freddie Mac. AIG holds

loss exceedances but found that the 10% (unconditional) quantile optimally balances the trade-off between sufficient numbers of nonzero observations in E_T^{-i} and sufficiently many extreme losses.




Figure 3. Full network graph for the system of the fifty-seven largest financial companies in the USA. For simplicity, arrows only mark risk spillover effects without referring to their respective size. Otherwise, arrows are as defined in Figure 1. Colors of the nodes mark insurers in black, broker dealers in the lightest gray, depositories in a darker gray and others in the darkest gray. A complete list of firms' acronyms is contained in Table I. The graphical allocation is obtained via the Fruchtermann–Reingold algorithm used to minimize the total length of all arrows.

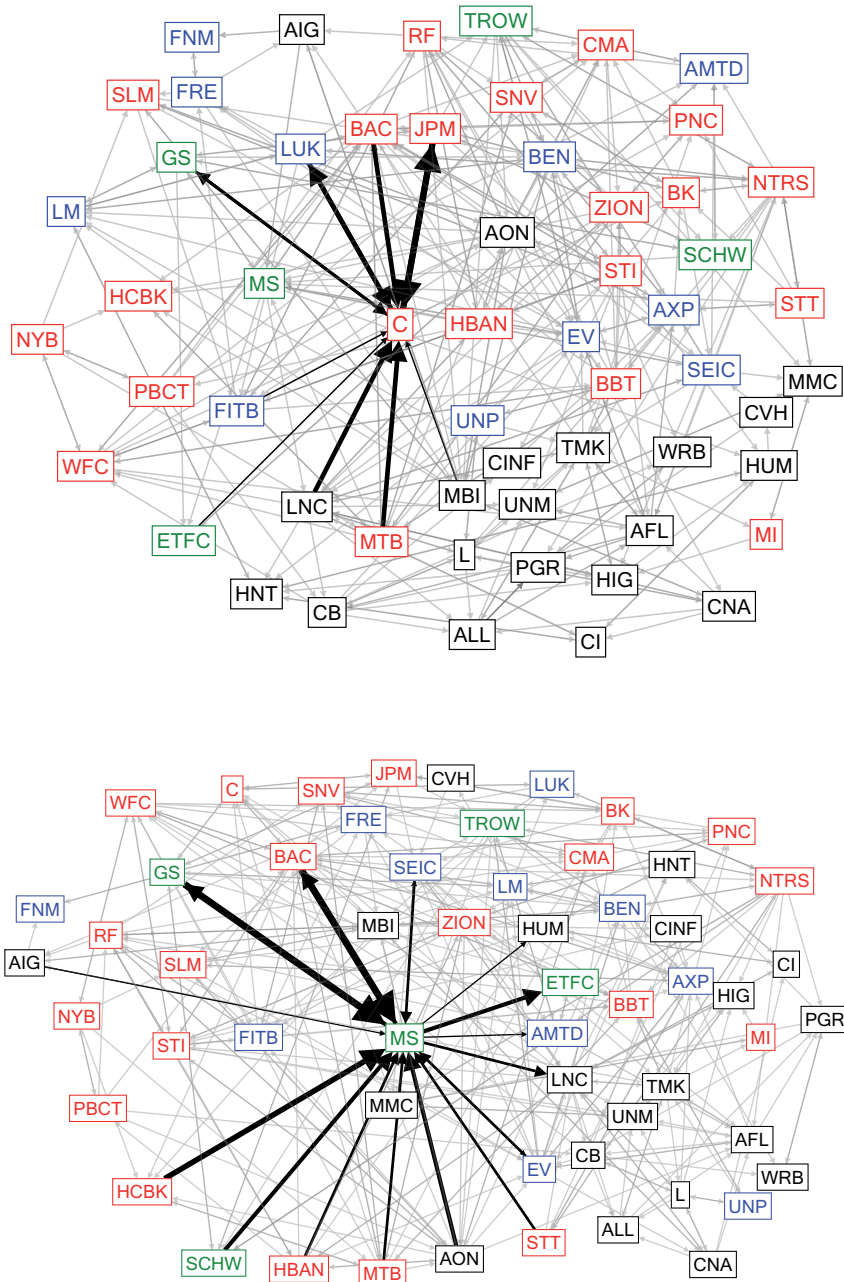


Figure 4. Full Network graphs of Citigroup (C) and Morgan Stanley (MS) highlighting risk drivers and risk recipients directly connected to the respective companies with bold arrows, according to the respective size of the effect. Arrows, colors, and acronyms are as in Figure 3. For simplicity, all other links only indicate spillover effects without referring to size. The list of firm acronyms is contained in Table I.




Figure 5. Network graph arranged according to industry groups, highlighting the industry-specific risk spillovers from depositories (top left), insurance companies (top right), broker dealers (bottom left), and others (bottom right). Arrows mark risk spillovers effects without referring to their respective sizes. Otherwise, arrows and colors are as in Figure 3. A complete list of firms’ acronyms is contained in Table I.

significant positions in mortgage backed securities and thus is closely connected to both Fannie Mae and Freddie Mac. Although the numbers of their relevant risk connections within the network are limited, such firms can nevertheless critically impact the overall system. In the 2008 financial crisis, the dependence between Freddie Mac and Fannie Mae and their interactions with AIG had severe systemic consequences.

Figure 5 indicates that it is not sufficient to focus exclusively on sector-specific subnetworks, as interconnectedness of institutions is widespread between industry sectors. We observe that tail risks of depositories,

insurance companies, and other firms are relatively evenly distributed among all industry groups. Depositories are the most strongly connected industry sector and exhibit their strongest connections among themselves. This contrasts with the other industries, where cross-firm connections within groups are less intense. In addition, broker dealers differ from other industry categories in that they display a much more concentrated risk outflow.

3.3 ROBUSTNESS

3.3.a *Network model validity*

The validity of the network identified is confirmed by four analyses: First, the significance of network effects in the individual VaR specifications are formally tested using a joint significance test of the individually selected loss exceedances E_t^{-i} in (2). We conduct this analysis based on a quantile regression version of the F -test for joint linear hypotheses developed by Koenker and Bassett (1982). Our results show that the selected tail risk spillovers are highly significant in all but a few cases. See Table III for an overview of all cross-effects. Detailed test results are available from the authors upon request.

Second, the importance of including other companies' loss exceedances as risk drivers for company i is further supported by comparing the (in-sample) forecast performance of our specifications with corresponding models of VaR^i , using macroeconomic variables only (as in Adrian and Brunnermeier (2011)). According to the employed backtests, specifications allowing for cross-firm dependencies have strong predictive ability and are significantly superior to models that ignore network linkages. Figure 6 shows the distributions of the backtesting p -values implied by both models. Hence, intercompany linkages add crucial explanatory power in VaR specifications.

Third, network effects remain important when we alter the set of economic state variables \mathbf{M} by adding asset pricing factors such as the three Fama-French factors and the momentum factor of Carhart (1997).¹⁸ We find that in the presence of network exceedances, these factors are deselected in all cases by the LASSO method and thus have no additional explanatory power. Hence, tails of asset returns are driven by factors other than the equity risk premium (associated with conditional means of returns).

¹⁸ The data are downloaded from the website of Kenneth French at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.




Figure 6. Boxplots of backtesting p -values indicating the in-sample model fit (i.e., testing the null hypothesis of formal statistical adequacy) of VaR specifications including macroeconomic regressors only (left) and VaR specifications resulting from the LASSO selection procedure (6) (right).

Finally, our results show that the most significant information about cross-company dependencies in tail risk is primarily contained in contemporaneous loss exceedances \mathbf{E}_t^{-i} . In contrast, alternative VaR specifications that utilize contemporaneous returns X_t^{-j} or lagged loss exceedances \mathbf{E}_{t-1}^{-i} yield significantly inferior backtest performance, with the regressors mostly found to be insignificant in joint F-tests.¹⁹

3.3.b Accuracy of the LASSO selection step

The firm-specific LASSO penalty parameter λ^i is a crucial coefficient under our approach, as it determines the denseness of the risk network and influences the selection of control variables when estimating systemic risk betas in Section 4. It is chosen in a data-driven way, optimizing a backtest criterion (see Sections 3.1 and 6). To validate this model selection step and to assess whether the procedure prevents overfitting, we analyze the consequences of increasing the LASSO penalty parameter. Note that higher values of λ^i lead to the selection of smaller models. If our procedure had a tendency to overfit

¹⁹ The corresponding results are available upon request and are omitted here for brevity.

the tails, the overall goodness of fit would increase for higher values of λ^i . We check for this by increasing all penalty parameters by 10% and 20%.²⁰ We show that, based on backtest performance, overall goodness of fit deteriorates substantially. This is demonstrated by the three boxplots and illustrations of individual p -values in Figure 7. For higher values of λ^i , the p -values decrease, and thus statistical support for the null hypothesis of a good model fit declines. Likewise, joint significance tests do not support the exclusion of additional regressors due to higher penalties. In particular, newly deselected regressors are mostly significant (jointly with the selected ones). This finding is confirmed by the Bayesian Information Criterion (BIC) for quantile models proposed by Lee, Noh, and Park (2013). As shown in Figure 7, the BIC is increasing, indicating a less favorable model when the penalty parameter is increased. These evaluations support our choice of penalization and indicate that there is no evidence of a tendency to overfit the tails.

3.3.c *Network characteristics*

In addition to providing a graphical illustration, standard network characteristics provide a more comprehensive picture of the interconnectedness and the role of each network node in the system. In Figure 8, we depict firms' pagerank coefficient (see Brin and Page (1998)), a measure that does not simply count links but empirically weights their importance in an iterative scheme.²¹ Confirming the visual impression based on Figure 3, the most connected firms are Lincoln National Corporation, AON, Bank of America, TD Ameritrade, and Morgan Stanley. The graph also illustrates our above finding that depositories tend to have somewhat stronger network effects than other industry groups. Insurance companies divide into a group of highly connected firms, such as Lincoln National Corp., AON and MBI, and a group of less connected companies, such as AIG, Humana Incorp., Unum Group (UNM), and Cincinnati Financial Corp.

The degree of firms' interconnectedness and the specific topology of the network allow for the identification of possible risk channels in the system. Pagerank coefficients, like other network metrics, however, can only be used

²⁰ Increasing the penalties beyond 20% is not advisable because, for some VaRs, no regressors are selected anymore.

²¹ The idea is to assign to each node (i.e., company) a weight that is increasing in the number of connections with other nodes and the relative importance thereof. The more connected a firm is, the greater is its importance and thus the greater is the importance of its neighbor. The computation of the pagerank coefficient can be viewed as an eigenvalue problem that can be solved iteratively. For more details, see Berkhin (2005).

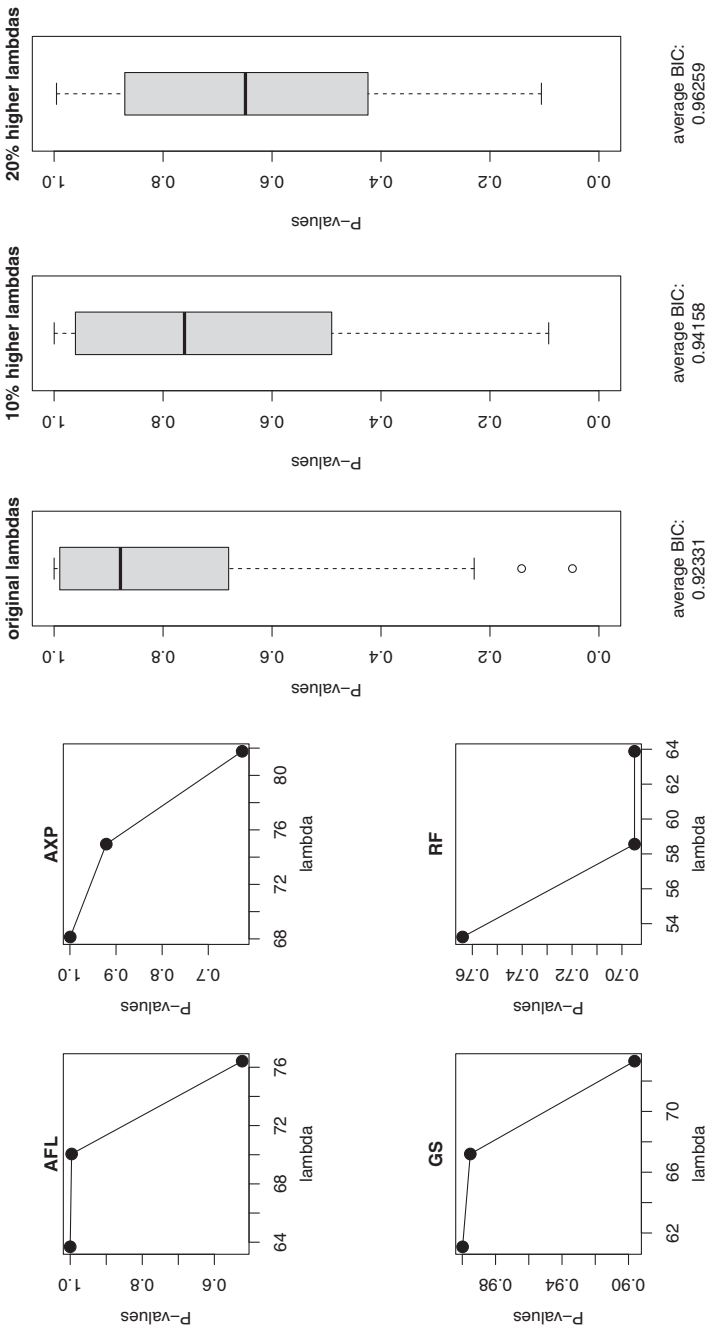


Figure 7. The left panel presents illustrative evolutionary paths of p -values from the VaR^i backtest described in Section 3.3.b. when individual-specific LASSO penalty parameters λ^i are increased by 10% and 20%. The respective leftmost p -value corresponds to the original choice. The right panel shows boxplots of all p -values obtained from backtesting all fifty-seven VaR time series. Higher p -values indicate better model fits. At the bottom of the right panel, average values of an additional goodness-of-fit measure, the BIC for quantiles, are reported. Lower values imply better model fits.




Figure 8. The upper figure displays pagerank coefficients based on the estimated tail risk network, computed as in Berkhin (2005), with institutions ordered by sector. Below, pagerank coefficients are plotted against realized systemic risk contributions measured as average realized systemic risk betas according to (8) for all companies classified as systemically relevant, according to Subsection 4.3, for the years 2000–08. The regression line shows only a small correlation between the pagerank coefficient and the realized systemic risk beta, a conclusion supported by the R^2 value of 0.0265. Colors and acronyms are as in Figure 3.

to assess the local impact and centrality of firms in the network but do not allow for a full quantitative assessment of the systemic relevance of a financial institution. To address the latter issue, we propose the concept of (realized) systemic risk beta, presented in the following section.

4. Quantifying Systemic Risk Contributions

4.1 MEASURING AND ESTIMATING SYSTEMIC RISK BETAS

In addition to valuable information on financial network structures, supervisory authorities seek an accurate but parsimonious measure of an institution’s systemic impact. We quantify the latter as the effect of a marginal change in the tail risk of firm i on the tail risk of the system, given the underlying network structure of the financial system. Similarly to a firm’s tail risk, as measured in Equation (1), system tail risk is measured as the Value-at-Risk $\text{VaR}_{p,t}^s$ of the system return X_t^s , conditional on $\text{VaR}_{q,t}^i$ and other control variables. We then define the systemic risk beta as the marginal effect of firm i ’s VaR on the system VaR given by

$$\frac{\partial \text{VaR}_{p,t}^s(\mathbf{V}_t^{(i)}, \text{VaR}_{q,t}^i)}{\partial \text{VaR}_{q,t}^i} = \beta_{p,q}^{s|i} \tag{7}$$

where $\mathbf{V}_t^{(i)}$ are firm-specific control variables.²² The systemic risk beta can be interpreted by analogy with an inverse asset pricing relationship in quantiles, where bank i ’s q -th return quantile drives the p -th quantile of the system, given network-specific effects and firm-specific and macroeconomic state variables. We classify the systemic relevance of institutions according to the statistical significance of $\beta_{p,q}^{s|i}$ at a given level and magnitude of their total effects

$$\bar{\beta}_{p,q}^{s|i} := \beta_{p,q}^{s|i} \text{VaR}_t^i \tag{8}$$

which we refer to as the realized systemic risk contribution. In contrast to the marginal systemic risk beta, the realized system risk beta captures the full partial effect of an increase in VaR_t^i on VaR_t^s and is thus cross-sectionally comparable across banks.

Producing unbiased estimates of a firm’s marginal effect $\beta_{p,q}^{s|i}$ requires accounting for i -specific control variables in (7). Consequently, for each company i , we estimate an individual quantile regression of VaR_t^s of the form

$$\text{VaR}_{p,t}^s = \mathbf{V}_t^{(i)'} \boldsymbol{\gamma}_p^s + \beta_{p,q}^{s|i} \text{VaR}_{q,t}^i \tag{9}$$

where the vector of regressors $\mathbf{V}_t^{(i)} = (1, \mathbf{M}'_{t-1}, \mathbf{VaR}_{q,t}^{(-i)'})'$ includes a constant effect, lagged macroeconomic state variables and the VaRs of all companies

²² We only study immediate effects of risk shocks of company i on the system and do not infer further steps. The latter would require additional dynamic modeling, which is beyond the scope of this analysis.

that are identified as risk drivers for firm i in Section 3. The resulting specifications are parsimonious, as they contain the minimum set of variables $\mathbf{V}_t^{(i)}$ that are necessary but sufficient to guarantee unbiased estimates of $\beta_{p,q}^{s|i}$.²³ Therefore, variables unrelated to VaR^i do not affect firm i 's systemic risk contribution and can be omitted.²⁴ We view this approach as a tractable alternative to a structural equilibrium model, which would require us to address involved specification issues. In addition, even if the latter model were correctly specified, it would yield imprecise estimates due to the high dimensionality and interconnectedness of the financial system, given limited data availability.

Systemic risk betas in (9) are allowed to be time-varying, accounting for periods of turbulence in which not only banks' risk exposures change, but also their marginal importance to the system may vary. We model potential time variations of $\beta^{s|i}$, using a linear model in lagged observable factors, \mathbf{Z}_{t-1}^i , which characterize a bank's propensity to experience financial distress,

$$\beta_{p,q,t}^{s|i} = \beta_{0,p,q}^{s|i} + \mathbf{Z}_{t-1}^i \eta_{p,q}^{s|i}, \quad (10)$$

where $\eta_{p,q}^{s|i}$ are parameters. Including \mathbf{Z}^i in lagged form renders systemic risk betas predictable, which is important for forward-looking monitoring and supervision of the financial system. Imposing linearity on $\beta_{p,q,t}^{s|i}$ in \mathbf{Z}_{t-1}^i yields stable estimates, given that these factors are updated only quarterly.²⁵

The time-varying systemic risk betas $\beta_{p,q,t}^{s|i}$ are estimated using (9) and (10), with the (unknown) VaR quantities VaR_t^i and $\mathbf{VaR}_{q,t}^{(-i)}$ replaced by the corresponding (post-LASSO) pre-estimates \widehat{VaR}_t^i and $\widehat{\mathbf{VaR}}_{q,t}^{(-i)}$ from (6).²⁶ Hence,

²³ Controlling for the relevant VaRs of other companies precludes simultaneity issues related to potential effects of VaR^s on VaR^i . Therefore, any remaining "reverse causality" can only stem from i -specific risk drivers that are not part of our sample but are constituents of the financial system portfolio. While this possibility cannot be completely neglected, given the composition of our sample, the remaining companies tend to be relatively small and unimportant. It thus appears quite unlikely that reverse causality may arise through this channel.

²⁴ See Angrist, Chernozhukov, and Fernández-Val (2006) for a simple Frisch-Waugh type argument in quantile regressions.

²⁵ More flexible functional forms (see e.g., Fan *et al.*, 2013) would substantially increase the computational burden and are not easily tractable, given the available data.

²⁶ Note that a direct one-step estimation is not feasible, as the individual parameters $\beta_{0,p,q}^{s|i}$ and $\eta_{p,q}^{s|i}$ could not be identified without the additional identification condition $Q_q(\varepsilon^i | \mathbf{W}_t^{(i)}) = 0$, implicitly bringing back the first-step estimation and model selection step. Moreover, inserting the linear individual VaR (2) into the linear system VaR model (9) yields

$$X_t^s = -\beta_{0,p,q}^{s|i} \widehat{\text{VaR}}_{q,t}^i - (\widehat{\text{VaR}}_{q,t}^i \cdot \mathbf{Z}_{t-1}^i)' \eta_{p,q}^{s|i} - \widehat{\mathbf{V}}_t^{(i)'} \boldsymbol{\gamma}_p^s + \varepsilon_t^s, \tag{11}$$

where $Q_p(\varepsilon_t^s | \widehat{\text{VaR}}_{q,t}^i, \widehat{\mathbf{V}}_t^{(i)}, \mathbf{Z}_{t-1}^i) = 0$. As in Section 3, estimates of all components of $\beta_{p,q,t}^{s|i}$ are obtained via a quantile regression minimizing

$$\frac{1}{T} \sum_{t=1}^T \rho_p \left(X_t^s + \mathbf{B}_t^{(i)'} \boldsymbol{\xi}^s \right) \tag{12}$$

in the unknown parameters $\boldsymbol{\xi}^s$, with $\mathbf{B}_t^{(i)} \equiv (\widehat{\text{VaR}}_t^i, \widehat{\text{VaR}}_t^i \cdot \mathbf{Z}_{t-1}^i', \widehat{\mathbf{B}}_t^{(i)'})'$. The resulting estimate of the full time-varying marginal effect $\hat{\beta}_{p,q}^{s|i}$ in (10) is then given by

$$\hat{\beta}_{p,q,t}^{s|i} = \hat{\beta}_{0,p,q}^{s|i} + \mathbf{Z}_{t-1}^i' \widehat{\eta}_{p,q}^{s|i}, \tag{13}$$

for given values \mathbf{Z}_{t-1}^i . Constant systemic risk betas occur as a special case under the restriction $\eta_{p,q}^{s|i} = 0$, yielding $\hat{\beta}_{p,q,t}^{s|i} = \hat{\beta}_{0,p,q}^{s|i} = \hat{\beta}_{p,q}^{s|i}$. An estimate of the realized beta (8) is obtained as $\hat{\beta}_{p,q,t}^{s|i} := \hat{\beta}_{p,q,t}^{s|i} \widehat{\text{VaR}}_t^i$.

The fact that certain regressors are not observed but only pre-estimated has crucial consequences for statistical inference. The quantile regression asymptotic standard errors obtained with commonly used software packages based on Koenker and Bassett (1978) are generally too small, as they do not account for the pre-step. In contrast to mean regressions, such two-step results are nonstandard in a quantile setting and are therefore provided in detail in (A.3) in the Appendix. To the best of our knowledge, these are new to the literature.

4.2 DETERMINING SYSTEMIC RELEVANCE

We determine a firm’s systemic relevance and the potential time variation thereof via formal statistical significance tests. As quantile versions of asymptotic *t*- or *F*-tests are not valid in finite samples, and simple direct bootstrap adaptations yield incorrect results for quantiles,²⁷ we perform

a full model of the system’s tail risk in observable characteristics. Model selection based on such a full model, however, is infeasible because correlation effects among the large number of regressors would produce unreliable results.

²⁷ Generally, asymptotic distributions often only provide a poor approximation of the true distribution of the (scaled) difference between the estimator and the true value when sample sizes are not sufficiently large. In the case of quantile regressions, this effect is even more pronounced, as valid estimates of the asymptotic variance have poor nonparametric rates and thus require even larger sample sizes to obtain the same precision.

finite-sample inference for a linear hypothesis H on $\beta_{p,q,t}^{s|i} \in \xi^s$ based on the test statistic

$$S_T = \min_{\xi^s \text{ under } H} \sum_{t=1}^T \rho_p(X_t^s + \mathbf{B}_t^{(i)'} \xi^s) - \min_{\xi^s \text{ unrestricted}} \sum_{t=1}^T \rho_p(X_t^s + \mathbf{B}_t^{(i)'} \xi^s), \quad (14)$$

where $\mathbf{B}_t^{(i)}$ is as defined above, and ξ^s denotes the corresponding parameter vector.²⁸ Note, however, that the asymptotic distribution of S_T involves unknown terms, so that a bootstrap procedure is needed. Conventional resampling techniques remain inconsistent for S_T , due to the nonsmooth objective function of the quantile regression. However, we can construct an adjusted “wild-type” bootstrap method that yields valid inference. This is described in detail in the Appendix (compare Chen *et al.*, 2008). For all tests below, we generally consider effects significant if p -values are below 10%.

We define a company as systemically relevant if an increase in its potential loss position, given economic state variables and i -specific risk inflows from other companies, entails significantly higher potential systemic loss. This requires that its systemic risk beta is significant and nonnegative.²⁹ Accordingly, we test for the joint significance of all components of $\beta_t^{s|i}$, using the hypothesis

$$\mathbf{H1} : \beta_0^{s|i} = \eta_{\text{MMM}}^{s|i} = \eta_{\text{SIZE}}^{s|i} = \eta_{\text{LEV}}^{s|i} = \eta_{\text{BM}}^{s|i} = \eta_{\text{VOL}}^{s|i} = 0.$$

Testing whether marginal effects on the system are indeed time-varying in firm-specific characteristics implies the joint hypothesis

$$\mathbf{H2} : \eta_{\text{MMM}}^{s|i} = \eta_{\text{SIZE}}^{s|i} = \eta_{\text{LEV}}^{s|i} = \eta_{\text{BM}}^{s|i} = \eta_{\text{VOL}}^{s|i} = 0.$$

If H2 is not rejected, we re-specify the systemic risk beta as a constant ($\beta_t^{s|i} = \beta^{s|i}$), re-estimate the model without interaction variables, and test the hypothesis $\mathbf{H3} : \beta^{s|i} = 0$.

²⁸ This test is an adaptation to the quantile setting of a method proposed by Chen *et al.* (2008) for median regressions.

²⁹ Because we do not impose a priori nonnegativity restrictions, systemic risk betas can become negative at certain points in time. In a few cases, we can attribute these effects to sudden time variations in one of the (interpolated) company-specific characteristics $\mathbf{Z}_{t-1}^{s|i}$, driving systemic risk betas temporarily into the negative region. These effects might be reduced by linking $\beta^{s|i}$ in (10) to (local) time averages of $\mathbf{Z}_{t-1}^{s|i}$, which would stabilize systemic risk betas at the cost of a potentially substantial loss of information. We see this as an alternative approach. However, we do not pursue it in the present context.

4.3 EMPIRICAL RESULTS AND ROBUSTNESS OF SYSTEMIC RISK BETAS

We estimate systemic risk betas according to the approach described in Section 4.1, based on $q = 0.05$ ³⁰ and using all firm-specific characteristics as potential drivers of time-variation in systemic risk betas, that is, $\mathbf{Z}_t^i = \mathbf{C}_t^i$. As a consequence, systemic risk contributions of two companies with the same exposure to macroeconomic risk factors and financial network spillovers may still differ due to their balance sheet structures.³¹

Table IV reports the p -values of the test described in Section 4.2, which is performed using the wild bootstrap procedure (illustrated above (A.9) in the Appendix), based on 2,000 re-samplings of the test statistic.³² We find that the majority of firms have a significant time-varying systemic risk beta. Conversely, for approximately 25% of the firms, systemic risk betas are insignificant. Table V ranks all systemically relevant companies for the period from 2000 to 2008, according to their average realized systemic risk contributions $\hat{\beta}^{s|i}$. The systemically most risky companies are found to be JP Morgan, American Express, Bank of America, and Citigroup. According to our network analysis above, these firms are strongly interconnected and thus should be closely monitored. Realized systemic risk betas, however, contain information on systemic relevance beyond a company's network interconnectedness. This is illustrated in Figure 8, which reveals only slightly positive dependencies between pagerank coefficients and realized systemic risk betas. With an R^2 of just 2%, this relationship, however, is not very strong. Hence, firms' interconnectedness is not sufficient to assess their systemic relevance.

As a first rough validity benchmark of our assessment, we compare our results with the outcomes of the Supervisory Capital Assessment Program (SCAP), conducted by the Federal Reserve in the spring of 2009, just after the end of our sample period. While we rely exclusively on publicly available market data, the Fed could draw on extensive nonpublic confidential balance sheet information that reveals credit and other risk interconnection channels among the nineteen largest US bank holding companies.³³ The

³⁰ As we set $p = q$, we suppress the quantile index.

³¹ Note that we keep the set of regressors M parsimonious, as described in Section 2.2 and justified in Subsection 3.3.a.

³² Because of multicollinearity effects, the interpretation of individual coefficients η might be misleading. Therefore, we refrain from reporting corresponding estimates.

³³ For details on SCAP, see Federal Reserve System (2009). To determine requested individual capital buffers under different market scenarios, the Fed's measure of systemic relevance uses proprietary information regarding risk interconnections. According to Huang,

Table IV. Classification of companies with significant and/or time-varying systemic risk betas according to p -values of the corresponding significance tests

In all cases, the test level is taken as 10% and firms are in alphabetical order within each category. P -values p_{H1} for the test on significance of systemic risk betas in the time period 2000–08 are depicted in column one (see Hypothesis H1 in Section 4.3). If $\hat{\beta}^{sli}$ is detected as being significant, a second test on time-variation of $\hat{\beta}^{sli}$ in firm-specific characteristics Z_t^i is performed yielding p -values p_{H2} (see Hypothesis H2 in Section 4.3). For firms with a significant but not a time-varying systemic risk beta (lower panel on the left, marked with stars), we reestimate the systemic risk beta without time-varying interaction terms and test again for its significance. These results (p_{H3}) are included in parentheses in the second column (see Hypothesis H3 in Section 4.3).

Companies with significant β^{sli}		
Name	p_{H1}	p_{H2}
AMERICAN EXPRESS	0.001	0.006
AMERICAN INTL.GP.	0.002	0.000
BANK OF AMERICA	0.002	0.001
CHARLES SCHWAB	0.019	0.013
CHUBB	0.017	0.015
CIGNA	0.001	0.013
CINCINNATI FINL.	0.010	0.004
CITIGROUP	0.026	0.066
COMERICA	0.016	0.020
FANNIE MAE	0.001	0.000
FIFTH THIRD BANCORP.	0.039	0.021
FRANKLIN RESOURCES	0.028	0.030
FREDDIE MAC	0.098	0.092
HARTFORD FINL.SVS.GP.	0.001	0.001
HUDSON CITY BANC.	0.043	0.035
HUNTINGTON BCSH.	0.010	0.011
LEGG MASON	0.026	0.060
LEUCADIA NATIONAL	0.041	0.016
LINCOLN NAT.	0.062	0.026
M & T BK.	0.033	0.021
MARSH & MCLENNAN	0.003	0.002
MARSHALL & ILSLEY	0.020	0.019
MORGAN STANLEY	0.041	0.095
PNC FINANCIAL SVS. GP	0.012	0.012
PROGRESSIVE OHIO	0.007	0.003
REGIONS FINANCIAL	0.034	0.029
STATE STREET	0.054	0.049
T ROWE PRICE GP.	0.090	0.076
TORCHMARK	0.002	0.001
UNION PACIFIC	0.040	0.035
UNUM GROUP	0.079	0.097
W R BERKLEY	0.007	0.037
WELLS FARGO & CO	0.015	0.027

(continued)

Table IV. (Continued)

Companies with significant β^{st}		
Name	p_{H1}	p_{H2}
ZIONS BANCORP.	0.095	0.100
AON*	0.063	0.192 (0.135)
E TRADE FINANCIAL*	0.072	0.160 (0.233)
JP MORGAN CHASE & CO.*	0.014	0.237 (0.047)
NY.CMTY.BANC.*	0.040	0.132 (0.088)
SEI INVESTMENTS*	0.014	0.115 (0.025)
TD AMERITRADE HOLDING*	0.049	0.131 (0.188)
Companies with insignificant β^{st}		
Name	p_{H1}	
AFLAC	0.220	
ALLSTATE	0.114	
BANK OF NEW YORK MELLON	0.199	
BB & T	0.120	
CNA FINANCIAL	0.410	
COVENTRY HEALTH CARE	0.257	
EATON VANCE NV.	0.276	
GOLDMAN SACHS GP.	0.667	
HEALTH NET	0.371	
HUMANA	0.189	
LOEWS	0.276	
MBIA	0.235	
NORTHERN TRUST	0.305	
PEOPLES UNITED FINANCIAL	0.105	
SLM	0.391	
SUNTRUST BANKS	0.213	
SYNOVUS FINL.	0.289	

financial institution with the most severe shortcomings with respect to capital buffer, according to the SCAP, was Bank of America, which ranks among our most highly systemically relevant companies, leading the ranking in June 2008 (Table VI, Panel B). In addition, we identify six of eight banks in our database that, according to the SCAP results, were threatened by financial distress under more adverse market conditions.³⁴ These results

Zhou, and Zhu (2010), the two network-based evaluations should be related, as companies with the highest detected systemic relevance in 2000–08 should carry the highest shares of hypothetical loss insurance premia. Consequently, they should face the highest requested increases in individual capital buffers in 2009.

³⁴ We detect Citigroup, FifthThird Bancorp, Morgan Stanley, PNC, Regions Financial, and Wells Fargo as systemically relevant. Due to a lack of data, we cannot include in our

Table V. Ranking of companies according to average realized systemic risk betas over the years 2000–08 Q3

Most systemic risk contributions are detected as time-varying in systemic risk betas—exceptions with constant $\hat{\beta}_{av}^{s|i}$ are marked by *. The underlying significance tests are performed as described in Table IV. The third column lists relevant risk drivers for the corresponding firm within the systemic tail risk network. They are determined through the LASSO selection technique (6) as “relevant” loss exceedances to be included in the corresponding company’s Var^i -regression.

Rank	Name	$\hat{\beta}_{av}^{s i} \cdot 10^2$	Influencing companies
1	JP MORGAN CHASE & CO	1.41*	BAC,BK,C,GS,PNC,SCHW
2	AMERICAN EXPRESS	1.22	AFL,BAC,BBT,BEN,CINF,EV,L,SEIC,SLM,STT, TROW
3	BANK OF AMERICA	1.01	AON,AXP,C,HBAN,LM,MS,MTB,PBCT,PNC, SEIC,STI,WFC
4	CITIGROUP	0.87	BAC,ETFC,FITB,GS,JPM,LNC,LUK,MBI,MTB
5	LEGG MASON	0.83	AON,BAC,BEN,CINF,EV,GS,HNT,MBI
6	REGIONS FINANCIAL	0.72	AMTD,AON,BBT,FITB,HBAN,PBCT,STI,ZION,
7	MARSHALL & ILSLEY	0.65	MMC,TMK
8	MARSH & MCLENNAN	0.63	MI,NTRS,PGR,SEIC,TROW,UNM
9	MORGAN STANLEY	0.62	AIG,AON,BAC,EV,GS,HBAN,HCBK,MTB,SCHW, SEIC,STT
10	AMERICAN INTL.GP.	0.61	FRE,MBI,RF,TMK
11	PROGRESSIVE OHIO	0.58	AFL,ALL,NTRS,WRB
12	STATE STREET	0.55	AXP,NTRS
13	ZIONS BANCORP.	0.51	BBT,CMA,HBAN,MTB,PNC,RF,STI,
14	FIFTH THIRD BANCORP.	0.49	AON,LUK,RF,SLM,SNV,STI,WFC,ZION
15	NY.CMTY.BANC.	0.49*	PBCT,WFC
16	PNC FINANCIAL SVS. GP	0.47	BAC,CMA,STT,TMK,WFC,ZION
17	FANNIE MAE	0.45	AIG,FRE
18	FRANKLIN RESOURCES	0.34	AON,AXP,BBT,EV,GS,LM,MBI,NTRS,SLM, SNV,TROW
19	CHARLES SCHWAB	0.33	AMTD,GS,JPM,NTRS,TROW
20	CHUBB	0.30	AFL,L,LNC,PBCT,PGR
21	WELLS FARGO & CO	0.28	BAC,BBT,CB,LNC,MTB,NYB,STI
22	FREDDIE MAC	0.19	BBT,EV,FITB,FNM,LUK
23	HARTFORD FINL.SVS.GP.	0.19	CB,L,LNC,MI,NTRS,TMK
24	CINCINNATI FINL.	0.16	CB,MBI,STI
25	TORCHMARK	0.12	AFL,ALL,BBT,HIG,LNC,NTRS,SEIC,UNM,UNP,
26	UNUM GROUP	0.04	AFL,ALL,L,LNC,MMC,STI

confirm the usefulness of our approach in detecting systemically risky companies and as a monitoring tool for supervisory agencies. For a more detailed validity study, see the next section below. Additionally,

analysis KeyCorp and GMAC, which have also been found to be financially distressed in critical macroeconomic environments.

Table VI. Rankings of relevant systemic risk contributions

Rankings are based on estimated realized systemic risk betas $\hat{\beta}_t^{sli}$ at two specific points in time. In addition, estimated systemic risk betas and VaRs are listed, illustrating the different sources of variation in $\hat{\beta}_t^{sli}$. Most systemic risk contributions are detected as being time-varying in systemic risk betas—exceptions with constant $\hat{\beta}_t^{sli}$ are marked by *. The underlying significance tests are performed as described in Table IV.

Panel A: End of March 2007 (before the beginning of the financial crisis)				
Rank	Name	$\hat{\beta}_{2007}^{sli} \cdot 10^2$	$\hat{\beta}_{2007}^{sli}$	$\widehat{\text{VaR}}_{2007}^i$
1	CITIGROUP	1.78	0.263	0.068
2	AMERICAN EXPRESS	1.35	0.387	0.035
3	BANK OF AMERICA	1.16	0.304	0.038
4	JP MORGAN CHASE & CO.	1.05*	0.265	0.040
5	MORGAN STANLEY	1.01	0.146	0.069
6	LEGG MASON	0.98	0.205	0.048
7	MARSH & MCLENNAN	0.83	0.222	0.037
8	REGIONS FINANCIAL	0.78	0.202	0.038
9	PNC FINANCIAL SVS. GP	0.77	0.248	0.031
10	CHUBB	0.74	0.240	0.031
11	AMERICAN INTL.GP.	0.61	0.143	0.043
12	FRANKLIN RESOURCES	0.60	0.143	0.042
13	STATE STREET	0.51	0.114	0.045
14	FIFTH THIRD BANCORP.	0.50	0.104	0.048
15	PROGRESSIVE OHIO	0.42	0.092	0.046
16	NY.CMTY.BANC.	0.41*	0.090	0.045
17	MARSHALL & ILSLEY	0.40	0.088	0.045
18	TORCHMARK	0.39	0.173	0.023
19	HARTFORD FINL.SVS.GP.	0.38	0.099	0.039
20	ZIONS BANCORP.	0.26	0.115	0.054
21	CHARLES SCHWAB	0.25	0.042	0.060
22	FREDDIE MAC	0.23	0.057	0.041
23	LEUCADIA NATIONAL	0.19	0.057	0.033
24	CINCINNATI FINL.	0.13	0.026	0.050
25	FANNIE MAE	0.09	0.019	0.049
26	UNUM GROUP	0.23	0.045	0.051
27	T ROWE PRICE GP.	0.06	0.014	0.043
28	LINCOLN NAT.	0.04	0.010	0.036

Panel B: End of June 2008 (during the financial crisis)				
Rank	Name	$\hat{\beta}_{2008}^{sli} \cdot 10^2$	$\hat{\beta}_{2008}^{sli}$	$\widehat{\text{VaR}}_{2008}^i$
1	BANK OF AMERICA	2.86	0.186	0.154
2	AMERICAN EXPRESS	2.78	0.278	0.100
3	WELLS FARGO & CO	2.51	0.186	0.135
4	MARSHALL & ILSLEY	2.31	0.516	0.045
5	JP MORGAN CHASE & CO.	2.22*	0.265	0.084
6	PROGRESSIVE OHIO	1.97	0.380	0.052
7	LEGG MASON	1.96	0.137	0.143

(continued)

Table VI. (Continued)

Panel B: End of June 2008 (during the financial crisis)				
Rank	Name	$\hat{\beta}_{2008}^{sl_i} \cdot 10^2$	$\hat{\beta}_{2008}^{sl_i}$	$\widehat{\text{VaR}}_{2008}^i$
8	REGIONS FINANCIAL	1.86	0.107	0.173
9	MARSH & MCLENNAN	1.76	0.471	0.037
10	STATE STREET	1.44	0.171	0.084
11	NY.CMTY.BANC.	1.12*	0.090	0.125
12	PNC FINANCIAL SVS. GP	1.09	0.153	0.071
13	CHUBB	1.07	0.176	0.061
14	TORCHMARK	1.00	0.177	0.057
15	CHARLES SCHWAB	0.91	0.149	0.060
16	CITIGROUP	0.90	0.072	0.124
17	MORGAN STANLEY	0.61	0.074	0.083
18	ZIONS BANCORP.	0.58	0.058	0.100
19	UNUM GROUP	0.34	0.033	0.104
20	UNION PACIFIC	0.27	0.047	0.056
21	HARTFORD FINL.SVS.GP.	0.24	0.012	0.201
22	FRANKLIN RESOURCES	0.17	0.026	0.064
23	T ROWE PRICE GP.	0.01	0.001	0.102

statistical robustness checks of an adapted realized systemic risk beta in a forecasting setting are provided in Hautsch, Schaumburg, and Schienle (2014).

To illustrate the time evolution of systemic risk betas, we show systemic risk rankings at two selected points in time: Table VI (Panel A) provides the systemic risk ranking for the last week of March 2007, which was a relatively “calm” time before the start of the financial crisis. Table VI (Panel B), on the other hand, shows the ranking at the end of June 2008, shortly before the collapse of Lehman Brothers. Comparing the precrisis and crisis rankings, we observe that systemic risk betas generally increased during the crisis, a trend that is particularly pronounced for American Express, Bank of America, JP Morgan, Regions Financial, and State Street. Exceptions are Citigroup and Morgan Stanley.

During the crisis, we detect Bank of America (BAC) as systemically most relevant. Among all systemically relevant companies, it is also the most interconnected firm, according to the pagerank coefficient in Figure 8, mutually influencing and influenced by other companies in the center of the network (Table III). Figure 9 shows that BAC’s systemic risk beta was relatively stable before the financial crisis but fell significantly following implementation of the Federal Reserve’s rescue packages. Its realized systemic risk contribution, however, strongly increased during the crisis,

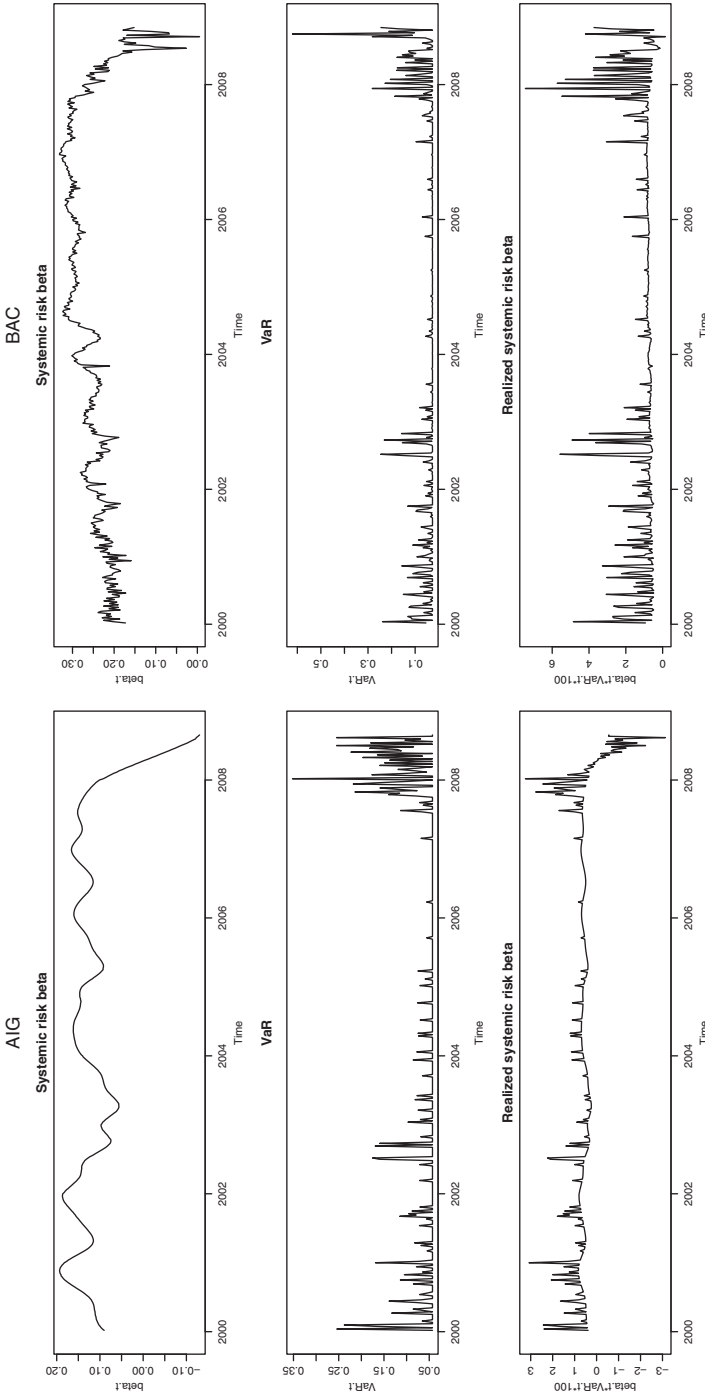


Figure 9. For each of the two institutions, AIG and Bank of America (BAC), the respective column comprises three time series panels, which depict, from top to bottom, the time-varying systemic risk beta $\hat{\beta}_t^{i,j}$, the time-varying VaR $\widehat{VaR}_t^{i,j}$, and the realized systemic risk beta $\hat{\beta}_t^{i,j} \widehat{VaR}_t^{i,j}$ of the firm.

a development that can mainly be attributed to network effects.³⁵ This is, for instance, in contrast to AIG, where network effects entering through increases in other firms' VaR's play a secondary role during the crisis period. AIG's systemic relevance, however, rapidly declined from the beginning of 2008 until its government bailout (Figure 9). Here, market data appear to have incorporated bailout information into the systemic risk beta well in advance, making both systemic risk betas and realized betas appear to be forward-looking.³⁶

By construction, realized systemic risk contributions vary over time through both β_i^{sl} and VaR_i . For selected companies, these effects are schematically depicted before and during the crisis in Figure 2 in the Introduction. As for BAC, in many cases, as shown in Table VI, we observe increases of realized systemic risk contributions that are mainly due to rising individual VaRs, while companies' marginal contributions to the system VaR generally remain unchanged (see, e.g., American Express). In most of these cases, the strong increase in VaR can mainly be attributed to tail risk spillovers in the network (see also Table II).

In several cases, increasing individual VaRs coincide with rising systemic risk betas. The most pronounced effect can be observed for Wells Fargo, which was not even identified as systemically relevant in 2007 but subsequently experienced a dramatic increase in both its systemic risk beta and its idiosyncratic tail risk, rendering it highly systemically risky in 2008. Other examples include State Street, Progressive Ohio, and Marshall & Isley. Here, direct sources of increasing systemic relevance can only be partially found in the network structure (see, e.g., State Street, which does not face significant risk spillovers from other companies but has high systemic relevance). For two central nodes in the network, Citigroup, and Morgan Stanley, however, declining systemic risk betas overcompensate for increasing VaRs, resulting in overall declining systemic relevance. Similarly to AIG, for these firms, network effects play a minor (direct) role.

The above results show that realized systemic risk contributions conveniently condense information on banks' systemic importance, although the underlying forces driving variations in banks' systemic relevance can be quite different. Therefore, simultaneously analyzing and monitoring (i) network effects, (ii) sensitivity to micro- and macroeconomic conditions,

³⁵ The detailed BAC results for the post-LASSO coefficients in Table 2 are omitted for the sake of brevity but are available upon request.

³⁶ For details on the USD 150 billion rescue packages from the Federal Reserve, see Schich (2009).

and (iii) time-variations in systemic risk betas provides a complete picture of companies’ specific roles in the network and thus builds a solid basis for supervisory authorities to monitor systemic risk.

5. Model Validation

5.1 A SIMPLISTIC BENCHMARK

In this subsection, we illustrate the advantage of our two-step quantile regression approach compared with a one-step estimation of a “global” model of the system VaR. Under the two-step approach, we model the system VaR as a function of all companies’ loss exceedances and the set of macroeconomic state variables³⁷, that is,

$$\text{VaR}_{p,t}^s = \alpha_1^s + \mathbf{E}_t^{(s)'} \alpha_2^s + \mathbf{M}_{t-1}^{(s)'} \alpha_3^s, \tag{15}$$

where $\mathbf{E}_t^{(s)}$ and $\mathbf{M}_{t-1}^{(s)}$ are sets of loss exceedances and macroeconomic indicators, respectively, selected by the LASSO method. Positive values of the unknown quantile-specific coefficients $(\alpha_1^s, \alpha_2^s, \alpha_3^s)'$ indicate the degree of systemic relevance of each firm. We use an adaptive version of the LASSO procedure in (15), employing regressor-specific weights in the penalty.³⁸ This makes the procedure comparable to the use of firm-specific LASSO penalties in our two-step procedure.

Figure 10 summarizes how the group of systemically relevant companies identified by the simplistic benchmark estimation compares with that determined by the two-step approach reported in Table V.³⁹ First, there is considerable overlap of companies—mostly large depositories and insurance companies (Group 1)—found to be systemically relevant under both

³⁷ LASSO selection in a global system VaR model, based on all institutions’ pre-estimated VaRs, would yield imprecise results due to the vast amount of pre-estimated regressors and inherent multicollinearity effects. We therefore do not include individual (pre-estimated) VaRs but loss exceedances.

³⁸ The adaptive LASSO criterion thus minimizes $\frac{1}{T} \sum_{t=1}^T \rho_q(X_t^s + \alpha_1 + \mathbf{E}_t' \alpha_2 + \mathbf{M}_{t-1}' \alpha_3) + \lambda \frac{\sqrt{q(1-q)}}{T} \sum_{k=1}^{65} w_k \hat{\sigma}_k |\alpha_k|$. The weights w_k are computed as inverses of the absolute values of coefficients from an unrestricted quantile regression, $\hat{\sigma}_k$ is as in (6), and λ is determined as in Section A.2, where c is chosen via the in-sample VaR backtest of Berkowitz, Christoffersen, and Pelletier (2011) (see Section A.3). For details on the adaptive LASSO, see Wu and Liu (2009).

³⁹ In the benchmark case, the change in the short-term interest rate (yield3m) was also used as a regressor, in addition to the selected exceedances. The detailed results with the coefficients obtained are available from the authors upon request.

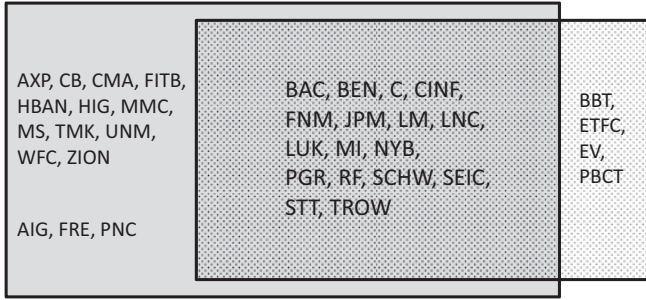


Figure 10. The schematic figure depicts companies classified as systemically relevant, according to our two-step network technique, in comparison with a simplistic one-step model, with exceedances based on LASSO for (15). Companies in the dotted area are selected by the simplistic model as systemically relevant, while firms in the gray area have a significant systemic impact in our network model, according to Table IV. Denote the overlay region as Group 1, which includes companies whose tail risks are determined as relevant to the system's risk in both settings. Group 2 comprises companies in the dotted but nongray area selected only by the simplistic model. Systemically relevant firms in the gray nondotted region can be classified as either Group 3, as they are deeply interconnected with other companies through more than six links, according to Table III (upper larger only gray set in the figure), or as Group 4, with few but crucial risk links, according to Table III (lower only gray set in the figure with three elements).

methods. In particular, 17 of 21 loss exceedances are selected under both approaches. Four remaining firms are identified as relevant only in the benchmark case (Group 2). These firms are relatively small companies that appear to be “overweighted” under the simplistic approach. The fact that they have been selected may indicate a spurious effect due to co-movements with others. The third group of companies comprises firms that are not identified as systemically relevant in the benchmark case but nevertheless have significant positive systemic risk betas. Almost all of these are deeply interconnected with other companies (Table III).⁴⁰ Hence, a one-step approach to determining the system VaR (15) may provide only a rough tool that can be used to gain a first impression of systemically relevant firms in a moderately interconnected system. However, as this approach cannot capture network linkages, it tends to systematically falsely reject systemic relevance of firms that gain importance mainly through their positions within the network. Conversely, it is likely to falsely attribute systemic relevance to firms with insignificant marginal effects when controlling for the network.

⁴⁰ We categorize a company as deeply connected if it has six or more incoming and/or outgoing risk links in Table III.

Table VII. Group ranking of systemic risk contributions for the precrisis period 2000 to mid 2007

The upper part, Group 1 (“high”), contains companies with significant average realized systemic risk betas in the highest quartile: $\hat{\beta}_{av}^{si} \cdot 100 \in [0.5, 1.3]$. Group 2 refers to the third quartile (“medium”) with $\hat{\beta}_{av}^{si} \cdot 100 \in [0.03, 0.49]$ and Group 3 to realized systemic risk betas lower than the median value (“small”), for which $\hat{\beta}_{av}^{si} \cdot 100 < 0.01$. Group 4 includes companies not determined to be systemically risky during the estimation period, that is, those with insignificant systemic risk betas. Case study companies are marked in bold.

Systemic risk contributions	Companies
Group 1 “high”	AIG, LEH , MS, JPM, GS, STT, CINF, LM, PBCT
Group 2 “medium”	FRE, ML , BAC, C, RF, AXP, PNC, CNA, TROW, NTRS
Group 3 “low”	FNM, WFC, EV, TMK, BBT, AFL, HUM, MI, CMA, BK, LNC, ALL, HNT, CB, CVH, SLM, ETFC
Group 4	AMTD, AON, BEN, CI, FITB, HBAN, HCBK, HIG, L, LUK, MBI, MMC, MTB, NYB, PGR, SCHW, SEIC, SNV, STI, UNM, UNP, WRB, ZION

5.2 CASE-STUDY: PRECRISIS PERIOD

The above study is based on data available over the entire period from the beginning of 2000 to the end of 2008. Consequently, institutions that defaulted or were taken over by other firms are not included. Nevertheless, to validate our findings, we perform a case study by reestimating the model for the time period of January 1, 2000, to June 30, 2007 and including the investment banks Lehman Brothers and Merrill Lynch.

The use of a shorter estimation period (and thus less data) renders a sharp ranking of companies less distinct and more difficult to interpret. Therefore, Table VII categorizes firms into groups according to quartiles of the distribution of realized systemic risk betas. The group of highest systemic importance consists of **AIG**, Lehman Brothers, Morgan Stanley, JP Morgan, and GS, among others. “Medium” systemic riskiness is observed for large depositories and investment banks, including Bank of America, Merrill Lynch, Citigroup, and Regions Financial and also for the mortgage company Freddie Mac.

In this case study, we focus particularly on four companies that were massively affected by the crisis: Lehman Brothers became insolvent on September 15, 2008 and was subsequently liquidated. Merrill Lynch announced a merger with Bank of America in September 2008, which was executed on January 1, 2009. Freddie Mac is closely connected to the second largest real estate financing company Fannie Mae and was placed under

Table VIII. Summary of estimation and test results for the four case study companies: loss exceedances influencing each company's VaR, the most important other VaRs influenced, joint significance tests on $\beta_t^{sli} = 0$ and estimated average systemic risk contributions as well as betas

^aTime-varying betas.

Estimation period: January 2000 to June 2007					
Name	Influenced by	Mainly influencing	Overall sign.	Average $\hat{\beta}_t^{sli} \cdot 100$	Average $\hat{\rho}_t^{sli}$
FRE	AON, BBT, EV, FITB, FNM, HUM, MBI	BBT, FNM	0.048	0.38	0.092 ^a
ML	AMTD, CB, CNA, HCBK, L, NYB, WRB	C	0.051	0.03	0.030 ^a
LEH	AMTD, AON, BEN, GS, JPM, LM, LUK, MI, MS	AIG, AXP, ETFC, JPM	0.041	0.79	0.176 ^a
AIG	ALL, C, CB, CNA, ETFC, HIG, LEH, LNC, MBI, MMC, SCHW, STT, TMK	AFL, C, CNA, HIG, HUM, MMC, UNM	0.026	0.73	0.210 ^a

conservatorship by the US government during the financial crisis. Finally, we investigate the systemic riskiness of AIG, which faced major distress during the crisis and whose bailout was very costly to tax payers. As shown in Table VII (with the specific companies marked in bold), all of these firms belong to the group of systemically relevant firms with high or mid-sized average systemic risk betas.

Table VIII summarizes the results of the network analysis of these four companies, using precrisis data only. We observe that most of these institutions were subject to loss spillovers from direct competitors. Observe, for example, the strong interconnectedness of AIG, the mutual link between Freddie Mac and Fannie Mae and the dependencies between Lehman Brothers and both MS and GS.

Figure 11 shows the time evolution of the realized betas of the four companies under investigation. The exemplary case of Merrill Lynch shows that, over a longer time horizon, the network-based idiosyncratic VaR gradually decreased, despite the firm's increasing systemic importance, with its realized risk beta rising by more than 100% between mid-2006 and mid-2007. Moreover, Figure 12 shows that the overall high systemic relevance of Lehmann and AIG can be attributed to very different time evolutions of their realized systemic risk betas well in advance of the crisis. While the systemic relevance of Lehman brothers grew almost monotonically toward the beginning of the crisis, the realized beta of AIG was already




Figure 11. Time evolution of systemic importance for all companies in the focus of the case study. The left column of the panel depicts quarterly averaged realized systemic risk betas of AIG, Freddie Mac (FRE), and Lehman Brothers (LEH) during the period immediately preceding the crisis. The right column shows quarterly averaged realized systemic risk betas of Merrill Lynch (ML) for the longer time period from 2004 onward in comparison with its VaR.

high around 2005. At this time, the company was already highly leveraged and was even downgraded. If we compare our results to the findings presented in Table IV (page 45 in the Appendix) of Brownlees and Engle (2012), according to their SRISK measure, they also find systemic relevance of LEH, FRE, and ML. In contrast to their results, however, our measure appears to incorporate important market information substantially more quickly, thus providing a better forward-looking monitoring tool. Similarly, the high systemic relevance of JPM before the crisis is detected by SRISK with a significant time delay.

Our findings clearly show that, in June 2007, all four companies were relevant to the stability of the US financial system. They indicate that bailouts during the crisis were justified for Freddie Mac (and the closely-tied Fannie Mae) and AIG. In addition, a failure of Merrill Lynch would have had harsh systemic consequences that could be prevented by its merger with Bank of America in 2008. Second, the increasing systemic importance of Lehman Brothers could have been monitored, and thus, the impact of its bankruptcy could have been anticipated in some degree. The direct




Figure 12. Time evolution of systemic importance in terms of quarterly realized systemic risk betas in 2004–07 for two companies, AIG and Lehman Brothers (LEH), which are among the most systemically risky companies. We depict quarterly averages, reflecting quarterly observations of balance sheet characteristics, smoothing the exceedance effects in the VaR's.

bidirectional linkage of Lehman Brothers to JP Morgan as well as the connections to MS and GS, which in turn are deeply interconnected, indicate a high risk of contagion in the event of Lehman's failure. Furthermore, our estimates show that Lehman's systemic risk contribution is only slightly lower than that of AIG, while it is substantially higher than that of, for example, Freddie Mac. Given these results, bailing out the latter but not the former is not necessarily justifiable from a systemic risk management point of view.

6. Conclusion

The global financial crisis of 2007–09 has demonstrated the need for an improved understanding of systemic risk. Particularly in situations of distress, it is the interconnectedness of financial companies that plays a

major role but challenges quantitative analysis and the construction of appropriate risk measures.

In this article, we propose a measure of firms' systemic relevance that accounts for dependence structures within the financial network, given market externalities. Our analysis allows us to statistically identify relevant channels of potential tail risk spillovers between firms, where such channels constitute the topology of the financial network. Based on these relevant company-specific risk drivers, we measure a firm's idiosyncratic tail risk by explicitly accounting for its interconnectedness with other institutions. Our measure of a company's systemic risk contribution quantifies the impact on the risk of distress of the system as a whole induced by an increase in the risk of an individual company in a network setting. Both measures exclusively rely on publicly observable balance sheet and market characteristics and can thus be used in prudent supervisory decisions in a stress test scenario.

Our empirical results show the interconnectedness of the US financial system and clearly mark channels of relevant potential risk spillovers. In particular, we classify companies into major risk producers, transmitters, or recipients within the system. Moreover, at any specific time, firms can be ranked according to their estimated contribution to systemic risk, given their roles and positions in the network. Monitoring companies' systemic relevance over time thus allows us to detect those firms that are most central to the stability of the system. In a case study, we highlight that our approach could have served as a solid basis for a sensible forward-looking monitoring tool before the start of the financial crisis in 2007.

Our approach is readily extendable in several directions. In particular, although the financial system is dominated by the USA, it truly is a global business with many firms operating internationally. Detecting inter- and intra-country risk connections and measuring firms' global systemic relevance should be straightforward under our proposed methodology. Moreover, whenever additional (firm-specific or market-wide) information becomes available, as, for example, when new information is reported to central banks, it can be directly incorporated into our measurement procedure. The data-driven selection step of relevant risk drivers then determines whether and how such information would increase the precision of results.

Appendix

ECONOMETRIC METHODOLOGY

Asymptotic results for two-step quantile estimation

Under the adaptive choice of the penalty parameter, as described in the text, the LASSO selection method is consistent with rate $O_P\left(\sqrt{\frac{K(i)}{T}} \log(\max(K, T))\right)$ and with high probability the coefficients selected from \mathbf{W} contain the true coefficients in finite samples. These results follow directly from Belloni and Chernozhukov (2011). Furthermore, VaR^i is consistently estimated, using the post-LASSO method described in the text, which reestimates the unrestricted model with $\mathbf{W}^{(i)}$. In particular, for all $q \in I$, with $I \in (0, 1)$ compact,

$$\widehat{\xi}_q^i - \xi_q^i \leq O_P\left(\sqrt{\frac{K(i)}{T}} \log(\max(K, T))\right), \tag{A.1}$$

as, in our setting, it is safe to assume that the number of wrongly selected components of \mathbf{W} is stochastically bounded by the number $K(i)$ of components of \mathbf{W} contained in the true model of VaR^i (see equation (2.16) in Belloni and Chernozhukov (2011)). Slightly abusing notation, we write $Y_T \leq O_P(r_T)$, with Y_T either $O_P(r_T)$ or even $o_P(r_T)$ for any random sequence Y_T and deterministic $r_T \rightarrow 0$. Note that, in general, for $T \rightarrow \infty$, both K and $K(i)$ might grow only very slowly in T , such that they can be treated nearly as constants, implying the standard oracle bound $O_P(\sqrt{\frac{\log(T)}{T}})$ in (A.1).

If the true model is selected, we find, for the asymptotic distribution of the individual VaR estimates for any $q \in [0, 1]$,⁴¹

$$\sqrt{\frac{1}{T}}(\widehat{\xi}_q^i - \xi_q^i)' \rightarrow N\left(0, \frac{q(1-q)}{g^2(G^{-1}(q))} \mathbb{E}[\mathbf{W}^{(i)} \mathbf{W}^{(i)'}]^{-1}\right), \tag{A.2}$$

where $g(G^{-1}(q))$ denotes the density of the corresponding error ε^i distribution at the q -th quantile. This result is standard (see Koenker and Bassett, 1978). For the second step estimates, we derive the asymptotic distribution analogously to the two-step median results in Powell (1983)

$$\sqrt{\frac{K(i)}{T}}\left((\widehat{\beta}_{0,p,q}^{s|i}, \widehat{\eta}_{p,q}^{s|i}, \widehat{\gamma}_p^s)' - (\beta_{0,p,q}^{s|i}, \eta_{p,q}^{s|i}, \gamma_p^s)'\right) \tag{A.3}$$

⁴¹ Required assumptions of Belloni and Chernozhukov (2011) and quantile analogies to Powell (1983) are fulfilled in our setting.

$$\rightarrow \mathcal{N}\left(0, Q^{-1}\mathbb{E}\left[\frac{p(1-p)}{f^2(F^{-1}(p))}\rho_p(\varepsilon_t^s) - \frac{p(1-p)}{g^2(G^{-1}(p))}\beta_{p,q}^{s|i'}\left(\rho_p(\varepsilon_t^i), \rho_p^y(\mathbf{Z}_{t-1}\varepsilon_t^i)\right)\right]\right), \tag{A.4}$$

where in the scalar factor, $f(F^{-1}(p))$ is the density of the corresponding error ε^s at the p -th quantile, the function ρ_p^y of a vector applies ρ_p to each of its components, and $\beta_{p,q}^{s|i} = (\beta_{0,p,q}^{s|i}, \eta_{p,q}^{s|i})$. The remaining main part Q in the variance is given by $Q = H'\mathbb{E}[\mathbf{A}\mathbf{A}']H$ with $\mathbf{A} = (\mathbf{W}^{(i)}, \text{vec}(\mathbf{Z}_{t-1} \cdot \mathbf{W}^{(i)}), \mathbf{VaR}^{(-i)})$. Denote by \mathbf{I} and $\mathbf{0}$ identity and null matrices, respectively, and by $\mathbf{1}$ a vector of ones of appropriate dimension. Then,

$$H' = \begin{pmatrix} \text{diag}(\xi_{q,2}^i) & \mathbf{0} & \dots \mathbf{0} \dots & \dots \mathbf{0} \dots \\ \mathbf{0} & \text{diag}(\xi_{q,1}^i) & \dots \mathbf{0} \dots & \dots \mathbf{0} \dots \\ \mathbf{0} & \mathbf{0} & \text{diag}(\text{vec}(\mathbf{1}_{d_z} \cdot \xi_q^{i'})) & \dots \mathbf{0} \dots \\ \mathbf{I} & \mathbf{0} & \dots \mathbf{0} \dots & \dots \mathbf{0} \dots \\ \mathbf{0} & \mathbf{0} & \dots \mathbf{0} \dots & \mathbf{I}_{d_{(-i)} \times d_{(-i)}} \end{pmatrix},$$

where d_Z is the dimension of Z , which is three in our application, $d_{(-i)}$ is the dimension of $\mathbf{VaR}_t^{(-i)}$, and coefficients $\xi_{q,2}^i$ are the components of ξ_q^i for regressors that appear both in the first and second step. Correspondingly, $\xi_{q,1}^i$ are coefficients of regressors that appear only in the first step of the individual VaR regression. Note that in the variance matrix, there is a distinction in γ for parts of \mathbf{V} that are also controls in VaR^i and $\mathbf{VaR}_t^{(-i)}$, which only appear in VaR^s .

Choice of the company-specific LASSO penalty parameter λ^i

We determine λ^i in a data-driven way following a bootstrap type procedure, as suggested by Belloni and Chernozhukov (2011):

Step 1. Take T iid draws from $\mathcal{U}[0, 1]$ independent of $\mathbf{W}_1, \dots, \mathbf{W}_T$ denoted as U_1, \dots, U_T . Conditional on observations of \mathbf{W} , calculate the corresponding value of the random variable,

$$\Lambda^i = T \max_{1 \leq k \leq K} \frac{1}{T} \left| \sum_{t=1}^T \frac{W_{t,k}(q - I(U_t \leq q))}{\hat{\sigma}_k \sqrt{q(1-q)}} \right|.$$

Step 2. Repeat step 1 for $B = 500$ times, generating the empirical distribution of Λ^i , conditional on \mathbf{W} through $\Lambda_1^i, \dots, \Lambda_B^i$. For the confidence

level $\alpha \leq 1/K$ in the selection, set

$$\lambda^i = c \cdot Q(\Lambda^i, 1 - \alpha | \mathbf{W}_t), \quad (\text{A.5})$$

where $Q(\Lambda^i, 1 - \alpha | \mathbf{W}_t)$ denotes the $(1 - \alpha)$ -quantile of Λ^i , given \mathbf{W}_t and where $c \leq 2$ is a constant.

The choice of α is a trade-off between a high confidence level and a corresponding high regularization bias from high penalty levels in (6). As in the simulation results in Belloni and Chernozhukov (2011), we choose $\alpha = 0.1$, which suffices to obtain optimal rates of the post-penalization estimators below. Finally, the parameter c is selected in a data-dependent way, such that the in-sample predictive ability of the resulting VaR specification is maximized. (Belloni and Chernozhukov, 2011 proceed in a similar way). The latter is evaluated in terms of its best backtesting performance, according to the procedure described right below.

Backtest for the model fit for VaRⁱ

As suggested by Berkowitz, Christoffersen, and Pelletier (2011), for each institution i , we measure VaR exceedances as $I_t^i \equiv I(X_t^i < -\text{VaR}_{q,t}^i)$. If the chosen model is correct, then,

$$\mathbb{E}[I_t^i | \Omega_t] = q, \quad (\text{A.6})$$

where Ω_t is the information set up to t . The VaR is estimated correctly, if, independently for each day of the covered period, the probability of exceeding the VaR is q . Similarly to Engle and Manganelli (2004); Kuuster, Mittnik, and Paolella (2006); Taylor (2008), we include a constant, three lagged values of I_t and the current VaR estimate in the information set Ω_t . Then condition (A.6) can be checked by estimating a logistic regression model

$$I_t^i = \alpha + \mathbf{A}_t' \theta + U_t,$$

with covariates $\mathbf{A}_t = (I_{t-1}^i, I_{t-2}^i, I_{t-3}^i, \widehat{\text{VaR}}_{t-1}^i)'$. Denote by \bar{I} the sample mean of the binary response I_t^i , and define $F_{\log}(\cdot)$ as the cumulative distribution function of the logistic distribution. Then, under the joint hypothesis

$$\mathbf{H}_0 := \alpha = q \quad \text{and} \quad \theta_1 = \dots = \theta_4 = 0,$$

the asymptotic distribution of the corresponding LR test statistic is

$$\text{LR} = -2(\ln \mathcal{L}_r - \ln \mathcal{L}_u) \overset{a}{\sim} \chi_5^2. \quad (\text{A.7})$$

Here, $\ln \mathcal{L}_u = \sum_{t=1}^n [I_t^i \ln F_{\log}(\alpha + \mathbf{A}'_t \theta) + (1 - I_t^i) \ln(1 - F_{\log}(\alpha + \mathbf{A}'_t \theta))]$ is the unrestricted log likelihood function, which, under \mathbf{H}_0 simplifies to $\ln \mathcal{L}_r = n\bar{I} \ln(q) + n(1 - \bar{I}) \ln(1 - q)$.

Bootstrap procedure for the joint significance test

The asymptotic distribution of the test statistic introduced in Subsection 4.1,

$$S_T = \min_{\xi^s \in \Omega_0} \sum_{t=1}^T \rho_p(X_t^s - \mathbf{B}'_t \xi^s) - \min_{\xi^s \in \mathbb{R}^{k_B}} \sum_{t=1}^T \rho_p(X_t^s - \mathbf{B}'_t \xi^s), \tag{A.8}$$

involves the probability density function of the underlying error terms and is not feasible. Furthermore, bootstrapping S_T would directly yield inconsistent results. Therefore, we resample from the adjusted statistic

$$S_T^* = \min_{\xi^s \in \Omega_0} \sum_{t=1}^T w_t \rho_p(X_t^s - \mathbf{B}'_t \xi^s) - \min_{\xi^s \in \mathbb{R}^{k_B}} \sum_{t=1}^T w_t \rho_p(X_t^s - \mathbf{B}'_t \xi^s) - \left(\sum_{t=1}^T w_t \rho_p(X_t^s - \mathbf{B}'_t \hat{\xi}_c^s) - \sum_{t=1}^T w_t \rho_p(X_t^s - \mathbf{B}'_t \hat{\xi}^s) \right), \tag{A.9}$$

where $\hat{\xi}_c^s$ denotes the constrained estimate of ξ^s , and $\{w_t\}$ is a sequence of standard exponentially distributed random variables, with both mean and variance equal to one. According to Chen *et al.* (2008), the empirical distribution of S_T^* provides a good approximation of the distribution of S_T . Thus, if the test statistic S_T exceeds some large quantile of the resampling distribution of S_T^* , the null hypothesis is rejected.

The proposed testing method does not require resampling of observations but is entirely based on the original sample. This provides significant gains in accuracy in the two-step regression setting relative to standard pairwise bootstrap techniques. A pre-analysis shows that this wild bootstrap type procedure is valid in the presented form, as any serial dependence in the data is sufficiently captured by the regressors in the reduced-form relation not requiring block-bootstrap techniques.⁴²

References

Acharya, V. V., Pedersen, L. H., Philippon, T., and Richardson, M. (2010) Measuring systemic risk, Technical report, MIT.

⁴² Pairwise block-bootstrapping yields block lengths of one, according to the standard procedure of Lahiri (2001). The results are available upon request.

- Acharya, V. V., Engle, R., and Richardson, M. (2012) Capital shortfall: a new approach to ranking and regulating systemic risks, *American Economic Review: Papers and proceedings* **102**, 59–64.
- Adrian, T. and Brunnermeier, M. K. (2011) Covar, Technical report, Princeton University.
- Allen, F. and Gale, D. (2000) Financial contagion, *Journal of Political Economy* **108**, 1–33.
- Angrist, J., Chernozhukov, V., and Fernández-Val, I. (2006) Quantile regression under misspecification, with an application to the US wage structure, *Econometrica* **74**, 539–563.
- Belloni, A. and Chernozhukov, V. (2011) l_1 -penalized quantile regression in high-dimensional sparse models, *The Annals of Statistics* **39**, 82–130.
- Berkhin, P. (2005) A survey on pagerank computing, *Internet Mathematics* **2**, 73–120.
- Berkowitz, J., Christoffersen, P., and Pelletier, D. (2011) Evaluating value-at-risk models with desk-level data, *Management Science* **57**, 2213–2227.
- Bernanke, B. (2009) Financial reform to address systemic risk, Technical report, Speech at the Council on Foreign Relations.
- Billio, M., Getmansky, M., Lo, A. W., and Pelizzon, L. (2012) Econometric measures of connectedness and systemic risk in the finance and insurance sectors, *Journal of Financial Economics* **104**, 535–559.
- Boss, M., Elsinger, H., Summer, M., and Thurner, S. (2004) Network topology of the interbank market, *Quantitative Finance* **4**, 677–684.
- Brin, S. and Page, L. (1998) The anatomy of a large-scale hypertextual web search engine, *Computer Networks and ISDN Systems* **3**, 107–117.
- Brownlees, C. T. and Engle, R. F. (2012) Volatility, correlation and tails for systemic risk measurement. Technical report, available at SSRN: <http://dx.doi.org/10.2139/ssrn.1611229>.
- Brunnermeier, M. (2009) Deciphering the liquidity and credit crunch 2007–08, *Journal of Economic Perspectives* **23**, 77–100.
- Carhart, M. M. (1997) On persistence in mutual fund performance, *The Journal of Finance* **52**, 57–82.
- Chen, K., Ying, Z., Zhang, H., and Zhao, L. (2008) Analysis of least absolute deviation, *Biometrika* **95**, 107–122.
- Cocco, J., Gomes, F., and Martins, N. (2009) Lending relationships in the interbank market, *Journal of Financial Intermediation* **18**, 24–48.
- Degryse, H. and Nguyen, G. (2007) Interbank exposures: an empirical examination of systemic risk in the Belgian banking system, *International Journal of Central Banking* **3**, 951–979.
- Diebold, F. X. and Yilmaz, K. (2012) Better to give than to receive: predictive directional measurement of volatility spillovers, *International Journal of Forecasting* **28**, 57–66.
- Diebold, F. X. and Yilmaz, K. (2014) On the network topology of variance decompositions: measuring the connectedness of financial firms, *Journal of Econometrics* **182**, 119–134.
- Elsinger, H., Lehar, A., and Summer, M. (2006) Risk assessment for banking systems, *Management Science* **9**, 1301–1314.
- Engle, R.F. and Granger, C. W. J. (1987) Co-integration and error correction: representation, estimation, and testing, *Econometrica* **55**, 251–276.
- Engle, R. F. and Manganelli, S. (2004) Caviar: conditional autoregressive value at risk by regression quantiles, *Journal of Business and Economic Statistics* **22**, 367–381.
- Fan, Y., Härdle, W., Wang, W., and Zhu, L. (2013) Composite quantile regression for the single-index model. Technical report, HU Berlin, CRC 649 working paper 2013-10.
- Federal Reserve System, Board of Governors (2009) The supervisory capital assessment program: overview of results, Technical report, FED.

- Freixas, X., Parigi, B., and Rochet, J. C. (2000) Systemic risk, interbank relations, and liquidity provision by the central bank, *Journal of Money, Credit, and Banking* **32**, 611–638.
- Friedman, N., Geiger, D., and Goldszmidt, M. (1997) Bayesian network classifiers, *Machine Learning* **29**, 131–163.
- Furfine, C. H. (2003) Interbank exposures: quantifying the risk of contagion, *Journal of Money, Credit and Banking* **35**, 111–128.
- Giesecke, K. and Kim, B. (2011) Systemic risk: what defaults are telling us, *Management science* **57**, 1387–1405.
- Hautsch, N., Schaumburg, J., and Schienle, M. (2014) Forecasting systemic impact in financial networks, *International Journal of Forecasting* **30**, 781–794.
- Huang, X., Zhou, H., and Zhu, H. (2010) Systemic risk contributions. Technical report, available at SSRN: <http://dx.doi.org/10.2139/ssrn.1650436>.
- Iyer, R. and Peydríó, J.-L. (2011) Interbank contagion at work: evidence from a natural experiment, *Review of Financial Studies* **24**, 1337–1377.
- Koenker, R. and Bassett, G. (1978) Regression quantiles, *Econometrica* **46**, 33–50.
- Koenker, R. and Bassett, G. (1982) Tests of linear hypotheses and l^1 estimation, *Econometrica* **50**, 1577–1584.
- Koopman, S. J., Lucas, A., and Schwaab, B. (2011) Modeling frailty correlated defaults using many macroeconomic covariates, *Journal of Econometrics* **162**, 312–325.
- Kuester, K., Mittnik, S., and Paolella, M. S. (2006) Value-at-risk prediction: a comparison of alternative strategies, *Journal of Financial Econometrics* **4**, 53–89.
- Kupiec, P. (1995) Techniques for verifying the accuracy of risk management models, *Journal of Derivatives* **3**, 73–84.
- Lahiri, S. N. (2001) Effects of block lengths on the validity of block resampling methods, *Probability Theory and Related Fields* **121**, 73–97.
- Lee, E. R., Noh, H., and Park, B. U. (2013) Model selection via Bayesian Information Criterion for quantile regression models, Technical report, Mannheim University.
- Leitner, Y. (2005) Financial networks: contagion, commitments and private sector bailouts, *Journal of Finance* **60**, 2925–2953.
- Pagan, A. R. and Wickens, M. R. (1989) A survey of some recent econometric methods, *The Economic Journal* **99**, 962–1025.
- Powell, J. L. (1983) The asymptotic normality of two-stage least absolute deviations estimators, *Econometrica* **51**, 1569–1575.
- Rajan, R. G. (2009) Too systemic to fail: consequences, causes and potential remedies, Technical report, Written statement to the Senate Banking Committee Hearings.
- Schich, S. (2009) Insurance companies and the financial crisis, *OECD Journal: Financial Market Trends* **15**, 123–151.
- Schwaab, B., Koopman, S. J., and Lucas, A. (2011) Systemic risk diagnostics: coincident indicators and early warning signals, Technical report, ECB Working Paper 1327.
- Segoviano, M. and Goodhart, C. (2009) Banking stability measures, Technical report, International Monetary Fund Working Paper 09/04.
- Taylor, J. W. (2008) Using exponentially weighted quantile regression to estimate value at risk and expected shortfall, *Journal of Financial Econometrics* **6**, 382–406.
- Tibshirani, R. (1996) Regression shrinkage and selection via the lasso, *Journal of the Royal Statistical Society, Series B* **58**, 267–288.
- Upper, C. and Worms, A. (2004) Estimating bilateral exposures in the German interbank market: is there a danger of contagion?, *European Economic Review* **48**, 827–849.

- White, H., Kim, T.-H., and Manganelli, S. (2010) Var for var: measuring systemic risk by multivariate regression quantiles. Technical report, available via Repec: <http://mpa.ub.uni-muenchen.de/35372>.
- Wu, Y. and Liu, Y. (2009) Variable selection in quantile regression, *Statistica Sinica* **19**, 801–817.
- Zhou, C. (2009) Are banks too big to fail? Technical report, De Nederlandsche Bank Working Paper 232-2009.