Pitfalls in the use of systemic risk measures*

Gunter Löffler
University of Ulm

Peter Raupach
Deutsche Bundesbank

July 6, 2015

Abstract

We examine pitfalls in the use of return-based measures of systemic risk contributions (SRCs). For both linear and non-linear return frameworks, assuming normal and heavy-tailed distributions, we identify non-exotic cases in which a change in a bank's systematic risk, idiosyncratic risk, size or contagiousness increases the risk of the system but lowers the measured SRC of the bank. Assessments based on estimated SRCs could thus produce false interpretations and incentives. We also identify potentially adverse side effects: A change in a bank's risk structure can make the measured SRC of its competitors increase more strongly than its own one.

Keywords: Systemic Risk; CoVaR; Marginal Expected Shortfall; Tail Risk
JEL classification: G21, G28.

1 Introduction

A measure of systemic risk aims to quantify how much an entity – be it a bank or hedge fund or sovereign – contributes to the vulnerability of the financial system. Recent years have seen a strong interest in refining such measures. Judging from the citation frequency, the two most influential concepts seem to be the CoVaR family of measures proposed by Adrian and Brunnermeier (2011) and the marginal expected shortfall of Acharya, Pedersen, Philippon, and Richardson (2012).[1] Originally intended for use in bank regulation, the literature now discusses these measures not only in conjunction with regulation but employs them for a variety of purposes: to examine whether systemic risk is priced (Meine, Supper, and Weiß (2015), Nucera, Schwaab, Koopman, and Lucas (2015)); to measure whether banks benefit from their too-big-to-fail status (Barth and Schnabel, 2013); to examine which funding channels or instruments are most important for systemic risk (López-Espinosa, Moreno, Rubia, and Valderrama (2012), Battaglia and Gallo (2013)); or to measure the contagion potential of sovereigns (Fong and Wong (2012)).

The literature on value at risk (VaR) has shown that the properties of risk measures and the consequences of choosing a specific measure for a particular purpose are not immediately obvious (cf. Artzner, Delbaen, Eber, and Heath (1999), Basak and Shapiro (2001)). With this

*Contact addresses: Gunter Löffler, Institute of Finance, University of Ulm, Helmholtzstrasse 18, 89069 Ulm, Germany. Phone: +49-731-5023597. e-mail: gunter.loeffler@uni-ulm.de. Peter Raupach, Research Center, Deutsche Bundesbank, Wilhelm-Epstein-Str. 14, 60431 Frankfurt am Main, Germany. Phone: +49-69-95668536. e-mail: peter.raupach@bundesbank.de (corresponding author). The views expressed in this paper represent the authors' personal opinions and do not necessarily reflect the views of the Deutsche Bundesbank or its staff. The paper replaces the Bundesbank Discussion Paper No 04/2013 "Robustness and informativeness of systemic risk measures".

[1] On 05/21/2015, Google Scholar showed 1051 citations for Adrian and Brunnermeier (2011) and 807 for Acharya et al. (2012).

paper, we want to contribute to a better understanding of systemic risk measures by pointing out pitfalls that can arise in typical applications. We examine whether the measures can give conflicting or misleading signals, e.g. indicate that a change in a bank's characteristics lowers its systemic risk contribution even though the change increases the risk of the system. Though we mainly interpret our results for a situation in which the measures are used in bank regulation, conclusions carry over to other applications. Wherever researchers build an analysis on measures of systemic risk, the reliability of their results will depend on the quality of the measures, and on how well their properties are understood.

The measures that we examine are the marginal expected shortfall (MES), the CoVaR, the exposure ΔCoVaR, and the beta. They have in common that stock returns are the key input for empirical measurement. The MES (Acharya et al. (2012)) is defined as an institution's average equity return on days in which the market return is below its 5% quantile. Originating in portfolio theory and suggested by its name (but not necessarily obvious from the definition), the MES is a marginal risk measure in that it equals the gradual change in a risk measure – the expected shortfall of the whole system's risk – when an institution gains weight in the system.

Risk contribution can also be understood as the change in systemic risk when an institution gets into distress rather than when it is added to the system. This is the intuition behind the ΔCoVaR measure introduced by Adrian and Brunnermeier (2011). It is the VaR of the system conditional on an institution being in distress, compared to the system's VaR conditional on the institution being in normal condition. Adrian and Brunnermeier also consider the VaR of an institution conditional on the system being in distress. The corresponding measure, called exposure ΔCoVaR, is more akin to portfolio theoretic concepts such as the MES. Several authors, e.g. Acharya et al. (2012) or Benoit, Colletaz, Hurlin, and Pérignon (2013), also analyze the beta, not necessarily as a candidate systemic risk measure but rather as a benchmark for the proposed measures. We include it for the same reason.

We start by examining within a linear market model framework how the measures respond to differences in systematic and idiosyncratic risk as well as in size. We find that the ΔCoVaR responds to idiosyncratic risk in an ambiguous way. When applied in regulation, the use of ΔCoVaR could create incentives for banks to increase idiosyncratic risk in order to lower their estimated systemic risk contribution. The way in which each of the four measures reacts to systematic risk (as measured through the exposure to a common factor) can cause similar problems, provided that an institution has a large weight in the system. With respect to size, finally, we find that the beta is susceptible to situations in which an increase in a bank's size lowers the estimated risk contribution.

For the sake of exposition, the presentation of results focuses on analytic derivations that assume a multivariate normal distribution for returns. Simulations with multivariate t-distributions as well as with a dynamic structural model show that the effects also appear in the presence of heavy tails and time-varying variances and sensitivities, and that they do not depend much on whether equity returns or asset returns are taken as a basis for the analysis.

Next, we examine a contagion framework in which negative shocks to one bank spill over to other banks. Even simple contagion structures can lead to a complex behavior of the four systemic risk measures. Some risk measures, notably the ΔCoVaR, have a tendency to assign a low systemic risk to infectious banks, others tend to do the opposite.

Our results are fundamental in the sense that we point out possible pitfalls that can arise in typical situations, with respect to key characteristics such as idiosyncratic risk or infectiousness. This should help in improving the use of existing measures, as well as in refining them. Deriving an exhausting set of axiomatic requirements for systemic risk measures and benchmarking existing proposals against them, though desirable, is beyond the scope of our paper. To illustrate why this is difficult to achieve, consider the case of contagion. Not only are there many definitions of contagion, as discussed by Pericoli and Sbracia (2003), it is also not evident which properties the measures should have. While it may seem intuitive to require that the systemic risk contribution of an infectious bank should, ceteris paribus, be larger than that of an infected bank, this is not obvious if the measure is used to identify risk control strategies. A regulator trying to increase the stability of a system could do so by either making banks more resilient that are likely to get infected, or by lowering the contagion intensity of contagious banks. Which option is preferred will depend on the costs of bringing about changes in infectiousness or resilience, the way in which banks respond to incentives, and other criteria applied by the regulator.

Closely related to our work are papers that explore the properties and limitations of systemic risk measures. Benoit et al. (2013) examine the similarity of risk rankings produced by different measures but do not examine whether the use of the measures can create unwanted incentives. Boucher, Kouontchou, and Maillet (2013) and Danielsson, James, Valenzuela, and Zer (2014) discuss measurement problems, from which we abstract in our analysis. Guntay and Kupiec (2014) suggest to separate systemic risk from systematic risk. While there may be situations in which it is useful to disentangle systematic risk from the effects of spillovers or interactions, we follow others (e.g. Bisias, Flood, Lo, and Valavanis (2012), and Allen and Carletti (2013)) and employ an inclusive definition of systemic risk because a systematic shock such as the bursting of a bubble can be sufficient to create jeopardizing system-wide losses.

Apart from the measures that we focus on because of their widespread application, there are several other return-based measures, many of which are related to MES and CoVaR; other branches of the literature employ holdings-based and network-based analyses of systemic risk. For an overview of the extensive literature, cf. Bisias et al. (2012), Benoit, Colliard, Hurlin, and Pérignon (2015), and Hüser (2015).

The remainder of the paper is structured as follows. In Section 2, we introduce the systemic risk measures studied in this paper. Section 3 discusses possible problems in a linear return setting, while Section 4 introduces contagion. Section 5 concludes.

2 Systemic risk measures studied in this paper

ΔCoVaR and exposure ΔCoVaR

Adrian and Brunnermeier (2011) suggest measures based on what they call CoVaR, which is implicitly defined through

$$\mathbf{P}\left(X^j \leq -CoVaR_\alpha^{j|C(X^i)} \middle| C\left(X^i\right) \right) = \alpha\,.$$

(To ease comparison with the other measures, we give the CoVaR the opposite sign of that in the original paper. All systemic risk measures in this paper will have the property that a higher value indicates a higher risk contribution.)

CoVaR is the value at risk (VaR) of object j conditional on event C happening to object i. Taking the event to be that i is at its VaR level, Adrian and Brunnermeier suggest to examine

$$\Delta CoVaR_\alpha^{j,i} = CoVaR_\alpha^{j|X^i=-VaR_\alpha^i} - CoVaR_\alpha^{j|X^i=Median^i}.$$

$\Delta CoVaR_\alpha^{j,i}$ measures the change in the α-VaR of j conditional on i moving from its median state to its own α-VaR. Adrian and Brunnermeier mostly examine the case in which j is given by the overall system, i.e. a market index or a collection of banks, and i is an individual institution; this is called ΔCoVaR throughout this paper. However, Adrian and Brunnermeier also consider the opposite direction in what they call exposure ΔCoVaR, which is defined through

$$\Delta CoVaR_\alpha^{j,system} = CoVaR_\alpha^{j|X^{system}=-VaR_\alpha^{system}} - CoVaR_\alpha^{j|X^{system}=Median^{system}}$$

$\Delta CoVaR_{\alpha}^{j,system}$ is the change in the VaR of portfolio j conditional on the system moving into distress. The measure is more akin to marginal expected shortfall and beta than $\Delta CoVaR_{\alpha}^{system,j}$.

Adrian and Brunnermeier estimate the CoVaR with a quantile regression over 25 years of weekly data, choosing a confidence level α of 1%. We abstract from estimation problems by deriving results through closed-form expression, or Monte Carlo simulations with a large number of observations.

Adrian and Brunnermeier suggest that, in the presence of time-varying risk, the precision of CoVaR estimates can be improved by conditioning the return-based estimates on current fundamental information. In most of our analysis, we consider static return frameworks. As we also abstract from estimation error, unconditional return-based estimates are optimal and fundamental information would not increase precision. The static framework also implies that a state-dependent modeling as in Adams, Füss, and Gropp (2014) would not enhance the informativeness of the CoVaR analysis.

In the base case, we do not model differences between asset returns – the use of which is advocated by Adrian and Brunnermeier – and equity returns. This is done for the sake of exposition, and seems justified given that for the short return horizons examined in the literature, the return distributions do not differ greatly except for their volatility. In a robustness check (Subsection 3.3), we show that conclusions are preserved insofar as all problems identified in the base case continue to exist when we use a dynamic structural credit risk model to differentiate between asset and equity returns, and when we introduce heavy tails and tail dependence into the return distributions of the base case model. However, further problematic effects appear that we do not observe in the base case.

Marginal expected shortfall (MES)

The marginal expected shortfall put forward by Acharya et al. (2012) is defined as

$$MES_i = -\mathbf{E}\left(R_i \,|R_S < Q_S^{\alpha}\right),$$

where R_i denotes the net equity return of institution i, R_S is the system return, and Q_S^{α} is the quantile of the system return on level α. Acharya et al. (2012) examine daily returns with a confidence level of 5%.

In the original work, the system return is proxied by the S&P 500, i.e. the authors include non-financial firms in the system. We deviate from this approach and follow Adrian and Brunnermeier, who empirically specify the system as consisting of financial institutions only. The system's scope matters in our analysis when we inspect the situation in which a single bank is very large compared to the system. Such a situation would be less likely if the system were taken to be the whole economy. Hence, our setup includes the perspective of a smaller country's financial system with a few big players (e.g. pre-crisis Iceland); their size is one of the parameters a systemic risk charge might react to in an undesired way.

Acharya et al. combine the MES measure with other information such as capital and size to calculate the systemic expected shortfall; see also Acharya, Engle, and Richardson (2012), who propose an extension with a focus on stochastic volatility. We stay with the MES since this additional information is beyond the scope of our analysis.[2]

[2] The systemic expected shortfall focuses on events that are actually critical to banks and likely to include default. Such events are typically not in the sample used for the estimation of a systemic risk measure. By contrast, the MES can directly be estimated and compared with the other measures considered in this paper.

Beta

Beta, the classic way to measure systematic risk, is defined as the regression coefficient in

$$R_i = a_i + beta_i R_S + u_i, \tag{1}$$

where R_i is the return of an individual bank and R_S is typically an index return. As with the other measures, we understand it as an index of banks. We write $beta_i$, not β_i, to highlight its role as a systemic risk measure, and reserve the Greek letter for the loading on a latent common factor.

While the difference between systemic risk (the danger of a breakdown of the financial system) and systematic risk (the exposure to common risk factors) is well acknowledged, a greater amount of the latter is likely to increase the former. For this reason, beta can also give an indication of systemic risk. Because of this and the widespread use of beta in finance, we examine its properties and compare them to those of other measures, similarly to Acharya et al. (2012) and Benoit et al. (2013).

Gauthier, Lehar, and Souissi (2012) test different capital allocation rules for their potential to improve system stability. Although their "component value-at-risk" allocation is also given the attribute "beta" in parentheses, it is more akin to the MES and the exposure ΔCoVaR at least in a linear setup with normal distributions (cf. Subsection 3.1 of Section 3). Also, and despite similar names, the concept of beta introduced here and that of tail beta introduced by Straetmans, Verschoor, and Wolff (2008) only have the idea in common that individual returns (or losses) are traced back to common factors. The tail beta is actually a conditional probability and designed to be invariant to marginal distributions. These are large conceptual and numerical differences to beta.

3 Systemic risk measures in the linear case

In this section, we use a linear factor model to examine whether the suggested measures for systemic risk fulfill elementary requirements with respect to a bank's choice of risk. We start with a linear combination of normally distributed factors, which we call the normal model. It allows analytical representations of the systemic risk measures, which we present first. Afterwards, we investigate in which way the measures depend on risk parameters. Assuming that the parameters are under control of the banks, the sensitivities reveal potential incentives that a systemic risk measure would hypothetically put on banks if applied as a systemic risk charge. Control over risk parameters will never be complete (and will have multiple other effects, e.g. on profitability), but clearly banks have more options to steer the risk of their business than average industrial firms.

As motivated in the introduction, we do not aim at judging systemic risk measures according to how useful they would be if systemic risk were regulated according to some optimum principle. We leave the way open in which a systemic risk charge or other regulatory measures are implemented but assume that banks desire to appear to be of low systemic importance, according to the systemic risk measure in place. Hence we inspect whether sensitivities to risk parameters have appropriate signs. Of course, other incentives may exist, in particular benefits from receiving a "too-big-to-fail" status. In cases in which such incentives are stronger than the opposite ones arising from the regulation of systemic risk, a change in the interpretation of the results may be required.

What matters for a bank is not only its own systemic risk but also the systemic risk of its competitors. If a bank's actions increase the systemic risk of competitors, the bank may have a desire to take such actions in order to gain a competitive advantage. We therefore consider different types of sensitivities: the impact of a bank's risk parameter on its own systemic risk measure as well as side effects on other banks.

We consider a banking system consisting of N banks. R_i, the return of bank i, is determined by the exposure to a common risk factor F and idiosyncratic risk ε_i. As explained above, we leave it open whether we have equity or asset returns in mind. Whether this neglect is justified will be discussed in Subsection 3.3 by means of a robustness check. Both factor returns and idiosyncratic components are assumed to be independent normal random variates. The size of bank i relative to the whole system is assumed to be w_i. Corresponding to the return type, size may be understood as market capitalization or total bank assets. We follow Adrian and Brunnermeier (2011) and take this system return to be the one that takes the role of a general market index, a choice that we also make in the estimation of MES and beta.

The system and its components are described through the following equations:

$$R_i = \beta_i F + \varepsilon_i,\ R_S = \sum_{i=1}^{N} w_i R_i \text{ with } F \sim N\left(\mu, \sigma_F^2\right),\ \varepsilon_i \sim N\left(0, \sigma_i^2\right),\ \sum_{i=1}^{N} w_i = 1, \tag{2}$$

where F and all ε_i are independent, β_i denotes the exposure to the common factor, and R_S is the return on the banking system index. We assume that all β_i are positive.

3.1 Analytic expressions for the risk measures

To calculate measures of systemic risk, we need to specify conditional distributions. Owing to the linearity of the system and the normality of the random variables, we can approach the problem in a linear regression framework. We start with an analysis of CoVaR measures. When we condition R_S on R_i, we study an orthogonal representation

$$R_S = c_i + d_i R_i + v_i,\quad R_i \perp v_i,$$

and obtain:

$$d_i = \frac{\text{cov}(R_S, R_i)}{\sigma^2(R_i)},\quad c_i = E(R_S) - d_i \mathbf{E}(R_i);\quad \sigma^2\left(v_i\right) = \sigma^2(R_S) - d_i^2 \sigma^2(R_i).$$

When we use ΔCoVaR to study how the system is affected by a distress of bank i, we obtain:

$$\begin{aligned}
\Delta CoVaR_\alpha^{S,i} &= -Q_\alpha\left(R_S \,|R_i = Q_\alpha\left(R_i\right)\right) + Q_\alpha\left(R_S \,|R_i = Q_{0.5}\left(R_i\right)\right) \\
&= -\left[c_i + d_i Q_\alpha\left(R_i\right) + Q_\alpha\left(v_i\right)\right] + \left[c_i + d_i Q_{0.5}\left(R_i\right) + Q_\alpha\left(v_i\right)\right] \\
&= -d_i\left[Q_\alpha\left(R_i\right) - Q_{0.5}\left(R_i\right)\right] = -d_i \sigma(R_i) \Phi^{-1}(\alpha),
\end{aligned}$$

with Φ denoting the standard normal cdf. Expanding d_i gives

$$\Delta CoVaR_\alpha^{S,i} = \frac{\text{cov}\left(R_S, R_i\right)}{\sigma\left(R_i\right)} \Phi^{-1}(1-\alpha). \tag{3}$$

We now turn to what Adrian and Brunnermeier (2011) call *exposure* Δ*CoVaR*. For this measure, R_i has to be conditioned on R_S rather than R_S on R_i. We therefore study

$$R_i = a_i + b_i R_S + u_i, \tag{4}$$

(which is also the definition equation (1) of the beta measure) to obtain:

$$b_i = \frac{\text{cov}(R_S, R_i)}{\sigma^2(R_S)},\quad a_i = \mathbf{E}(R_i) - b_i \mathbf{E}(R_S);\quad \sigma^2\left(u_i\right) = \sigma^2(R_i) - b_i^2 \sigma^2(R_S). \tag{5}$$

When we use the exposure ΔCoVaR to study how bank i is affected by the system, similar calculations as for the ΔCoVaR give:

$$\Delta CoVaR_{\alpha}^{i,S} \quad = \quad \frac{\text{cov}\,(R_S, R_i)}{\sigma\,(R_S)}\Phi^{-1}(1-\alpha) \tag{6}$$

The next measure considered is the marginal expected shortfall $MES_i = -\mathbf{E}\,(R_i\,|R_S < Q_S^{\alpha})$. As in exposure ΔCoVaR, a bank's return is conditioned on the system return. We therefore start from (4) to derive:

$$\begin{aligned} MES_i &= -\mathbf{E}\,(R_i\,|R_S < Q_S^{\alpha}) = -\mathbf{E}\,(a_i + b_i R_S + u_i\,|R_S < Q_S^{\alpha}) \\ &= -a_i - b_i \mathbf{E}\,(R_S\,|R_S < Q_S^{\alpha}) - \mathbf{E}\,(u_i\,|R_S < Q_S^{\alpha})\,. \end{aligned}$$

By construction of an OLS regression, u_i and R_S are uncorrelated. Because their joint distribution is multivariate normal (both are linear images of independent normals), they are independent, so that $\mathbf{E}\,(u_i\,|R_S < Q_S^{\alpha})$ is zero. We therefore obtain $MES_i = -a_i - b_i\,\mathbf{E}\,(R_S\,|R_S < Q_S^{\alpha})$ and, expanding a_i,

$$\begin{aligned} &MES_i = -\mathbf{E}\,(R_i) - b_i\,\mathbf{E}\,(R_S - \mathbf{E}\,(R_S)\,|R_S < Q_S^{\alpha}) \\ &\quad = \quad -\beta_i\mu - b_i\,\sigma\,(R_S)\,\mathbf{E}\left(Z\,\middle|Z < \Phi^{-1}\,(\alpha)\right), \quad Z \sim N\,(0,1) \\ &\quad = \quad -\beta_i\mu - \frac{\text{cov}\,(R_S, R_i)}{\sigma\,(R_S)}\mathbf{E}\left(Z\,\middle|Z < \Phi^{-1}\,(\alpha)\right) = -\beta_i\mu + \frac{\text{cov}\,(R_S, R_i)}{\sigma\,(R_S)}\frac{\phi\left(\Phi^{-1}\,(\alpha)\right)}{\alpha} \end{aligned} \tag{7}$$

The last transform is a familiar result for truncated normal distributions. Comparing (6) and (7), MES and exposure ΔCoVaR turn out to be essentially substitutes in a linear model of normals, in that they differ only by a constant factor and a shift that is typically small on a short-term horizon.

Beta has a very simple form in the linear setup. Comparing (1) with (4) and applying (5), it is $beta_i = \text{cov}\,(R_S, R_i)\,/\sigma^2\,(R_S)$. Note that the weighted average of all betas equals 1 by construction, which makes a difference in our sensitivity analyses where we let a single β_i increase while keeping the others constant. This is impossible for the $beta_i$, as there is always a compensating change in the $beta_j$ of other banks. The difference between $beta_i$ and β_i is therefore not just a notational one, besides the influence of idiosyncratic risks on $beta_i$.

The following formulas summarize the systemic risk measures in our linear model:

$$\Delta CoVaR_{\alpha}^{S|i} = \frac{\text{cov}\,(R_S, R_i)}{\sigma\,(R_i)}\Phi^{-1}(1-\alpha)\,; \qquad \Delta CoVaR_{\alpha}^{i|S} = \frac{\text{cov}\,(R_S, R_i)}{\sigma\,(R_S)}\,\Phi^{-1}(1-\alpha); \tag{8}$$

$$MES_i = -\beta_i\mu + \frac{\text{cov}\,(R_S, R_i)}{\sigma\,(R_S)}\,\frac{\phi\left(\Phi^{-1}\,(\alpha)\right)}{\alpha}\,; \qquad beta_i = \frac{\text{cov}\,(R_S, R_i)}{\sigma^2\,(R_S)}\,. \tag{9}$$

The covariance between the individual and the system return is common and central to all four measures; it is then scaled down by a specific variation measure either of individual or system returns. Only the MES deviates from this pattern by an additive component, of which we however show in Appendix A.7 that it is usually of low influence.

3.2 Sensitivities to risk parameters

With the analytic expressions at hand, we can now investigate the sensitivities of systemic risk measures to changes in a bank's risk parameters. We study effects of bank-specific idiosyncratic

(σ_i) and systematic risk (β_i), but also of size[3] (w_i). Size is not literally a risk parameter but has impact on the risk of the system, even in a simple setup as ours, where it only affects the weights of individual returns in the average return R_S. As argued above, two kinds of sensitivity appear to be relevant.

First, there is the effect that a bank's parameter has on its own systemic risk charge; we call it the *direct* effect and measure it by the partial derivative.

Second, there is also a potential side effect on other banks' systemic risk charges. A bank could benefit if its relative risk ranking compared to other banks improves. We measure this side effect through the change in the ratio of a bank's systemic risk to the systemic risk of another, representative bank. We call this sensitivity the *relative effect.* Summing up, for a parameter p_i applying to bank i and one of the systemic risk measures SRM considered in this paper, we calculate

$$\frac{\partial\, SRM_i}{\partial p_i} \text{ (direct effect) and } \frac{\partial}{\partial p_i}\left[\frac{SRM_i}{SRM_j}\right],\ j \neq i,\ \text{(relative effect)}.$$

We focus on two basic properties the sensitivities should fulfill in order to avoid unwanted incentives and interpretations. If a change in the parameter of a bank increases the risk of the whole system, it should also increase the systemic risk measure of the bank. In addition, it should not lower the systemic risk relative to other banks. To give the two properties a precise meaning, we must specify what we understand by the risk of the system. We define it as $\sigma(R_S)$ but note that in our multivariate normal setup, $VaR_\alpha(R_S)$ or expected shortfall $ES_\alpha(R_S)$ are almost synonym to the standard deviation, as regards monotonicity.[4] Both for idiosyncratic and systematic risk, there is no ambiguity about what sensitivity is appropriate. If β_i or σ_i rises, ceteris paribus, this always increases $\sigma(R_S)$, so that bank i should be assigned more systemic risk. Positive partial derivatives to β_i and σ_i are therefore appropriate.

By contrast, the case of size is ambiguous. If a bank's weight in the system increases, this can make $\sigma(R_S)$ rise or fall. For instance, when a bank with low β_i and moderate idiosyncratic risk gains weight, $\sigma(R_S)$ can fall because of the low β_i even though idiosyncratic risks become less diversified. Assigning such a growing bank less systemic risk would then be appropriate, at least as long as the role of size is limited to its impact on the joint distribution of returns.[5]

Evaluating all the sensitivities means analyzing the signs taken on by the elements of an 8×3 Jacobian and how they depend on other parameters. The results are summarized in Table 1. Since the sign of a sensitivity can depend on the parameterization, we provide the following information:

- In Panel A, we report the range of signs. A superscript ^{n} marks cases where the stated range applies under normal conditions. The single instance of "+/(–)" is discussed below.
- In Panel B, we report the signs of partial derivatives for a base case parameterization where parameters are set to $\mathbf{E}(F) = 0.05$, $\sigma_F = 0.2$, $\sigma_i = 0.2$, which are expressed on an annual basis; to translate them to daily returns, we divide $\mathbf{E}(F)$ by 260 and the standard deviations by $\sqrt{260}$. Further parameters are: $N = 50$, $\beta_i = 1$, $w_i = 1/50$ for all i; quantile level $\alpha = 0.01$ for the CoVaR measures and 0.05 for the MES.

[3] A change in size *ceteris paribus* is understood as a *dollar* change (of total assets or market capitalization, de-pending on the return used), while the dollar value of the other banks remains constant. When the weight w_i of a bank increases, the other banks' weights are therefore assumed to shrink proportionally, keeping the sum of all weights at 1.

[4] An opposite situation in which a parameter shift lets $\sigma(R_S)$ rise while $VaR_\alpha(R_S) = -Q_\alpha(R_S)$ shrinks would require the VaR to be negative. Even for a particularly moderate tail (10%), low system volatility (10%), and a long risk horizon (1 year), the annual drift of R_S would have to exceed 12.8% to make the system VaR negative.

[5] Changing size may have additional effects, which, however, cannot be analyzed in our simple linear model. For example, an increase in size could lower system stability because a larger banking sector makes it more likely that systemic threats cannot be contained through government intervention.

- In Panel C, we report signs for a system with a dominant bank of high systematic risk. Parameters are: $\beta_1 = 1.5$, $w_1 = 0.3$, $w_j = 0.7/49$ for $j > 1$; other parameters are as in the base case.

Overall, Table 1 provides a good impression, in that most of the effects (direct and relative) are either always positive or positive under normal conditions; only very implausible parameters constellations, specified in the appendix, would generate the opposite sign. There is one case marked by "+/(−)" in which the partial derivative can be negative under normal conditions, but only when the partial derivative to $\sigma(R_S)$ is negative, too. Hence, this case is not problematic within our framework. There are some exceptions, however, which we now look at in more detail.

If a bank increases its idiosyncratic risk, there are two opposing effects on ΔCoVaR: inspecting the fraction in the representation of $\Delta CoVaR_\alpha^{S|i}$ found in (8), the numerator $\text{cov}(R_S, R_i)$ increases in idiosyncratic risk since the bank is part of the system. The relevant addend to the covariance is given by $w_i\sigma_i^2$ so that the effect on the numerator becomes smaller, the lower the weight of the bank within the system. The denominator, which equals $\sqrt{\beta_i^2\sigma_F^2 + \sigma_i^2}$, is unaffected by the bank's weight, and the dependency on σ_i has a substantial linear part. Taken together, assuming moderate values for w_i and σ_i, the denominator's sensitivity in σ_i will be stronger than that of the numerator so that the ΔCoVaR can shrink while idiosyncratic risk rises. A systemic risk charge based on ΔCoVaR would therefore create an incentive for banks to increase their idiosyncratic risk, which would increase the volatility of the system return.

The partial derivative can be positive if the banks' weight or the idiosyncratic risk are large enough. To see which case probably prevails in practice we analyze the partial derivative of ΔCoVaR more deeply in the appendix.

In Figure 1 we pick a special case where all β_i equal 1. Then,

$$\frac{\partial CoVaR_\alpha^{S,i}}{\partial \sigma_i} \propto \frac{w_i}{1-w_i} - \left(1 + \frac{\sigma_i^2}{\sigma_F^2}\right)^{-1}, \tag{10}$$

which gives rise to setting the ratio σ_i/σ_F on the horizontal axis of Figure 1 and w_i on the vertical.

Points where the partial derivative is negative are displayed in gray. The area marks cases in which a bank striving for a low systemic risk measure might be tempted to increase its idiosyncratic risk. The area seems to cover most of the practically relevant cases. For instance, if the weight of a bank in the system is below 10 percent, the idiosyncratic risk can be up to 2.8 larger than its exposure to the systematic factor while the bank still has an incentive to increase idiosyncratic risk. Even if the bank makes up one third of the system, it still may be tempted to increase σ_i as long as this is not larger than σ_F.

To give an idea of the strength of the effect, we provide an example in Figure 2 where the ΔCoVaR is plotted as a function of σ_i (on an annual basis) between 0 and one half. Each line refers to a certain β_i; all other banks' β_j are held constant at 1 and their idiosyncratic risk at 0.2 (p.a.). All banks have the same weight; further parameters are given in the figure's notes. The upper graph confirms the negative dependency on σ_i. Bank i would benefit from large idiosyncratic risk most if its β_i is low. For the example of $\beta_i = 0.5$, assume the bank raises σ_i from 0.15 to 0.25. It could lower its ΔCoVaR from 0.0162% to 0.0113%, which is a reduction by 30.4 percent. For $\beta_i = 1.5$, where the sensitivity is weaker, the reduction would amount to 13 percent. The middle graph shows that idiosyncratic risk has no side effect on the ΔCoVaR of other banks (which also results from the formula) so that the lower graph, showing the relation between the ΔCoVaR of two banks, is just a duplicate of the upper.

However, direct and relative effects can also have opposite directions, as we show using the example of how ΔCoVaR reacts to systematic risk. In Figure 3, β_i varies between 0.5 and 2. The

Table 1: Effect of risk parameters on systemic risk measures.
We analyze sensitivities of systemic risk measures to certain risk parameters in a linear setting. Returns are described through

$$R_i = \beta_i F + \varepsilon_i,\ R_S = \sum_{j=1}^{N} w_j R_j \text{ with } F \sim N\left(\mu, \sigma_F^2\right),\ \varepsilon_i \sim N\left(0, \sigma_i^2\right).$$

All ε_i and F are independent. A direct effect of a parameter is understood as the partial derivative of a systemic risk measure. A relative effect refers to the ratio between the systemic risk measures of two banks. It is the ratio's partial derivative to a parameter of the bank in the numerator, e.g. $\partial\left(MES_i/MES_j\right)/\partial\sigma_i$.
Panel A presents possible signs of the derivatives. A superscript ^{n} marks cases where the sign applies under normal conditions. Only very implausible parameters constellations, specified in the appendix, would generate the opposite sign. In the single instance of "+/(−)", the partial derivative can be negative under normal conditions, but only when the partial derivative to $\sigma(R_S)$ is negative, too.
Panel B reports the partial derivatives' signs for the base case where parameters (p.a. for drift and volatility) are set to $N = 50$, $\mathbf{E}(F) = 0.05$, $\sigma_F = 0.2$, $\sigma_j = 0.2$, $\beta_j = 1$, $w_j = 1/50$ for all j; quantile level $\alpha = 0.01$ for the CoVaR measures and 0.05 for the MES.
Panel C reports signs for a system with a dominant bank of high systematic risk. Parameters are: $\beta_i = 1.5$, $w_i = 0.3$, $w_j = 0.7/49$ for $j \neq i$; other parameters as in the base case.

Parameter	Effect type	ΔCoVaR	Exp. ΔCoVaR	MES	Beta
Panel A: Range of the sign of partial derivative					
idiosyncratic risk σ_i	direct	+/−	+	+	+
	relative	+/−	+	+	+
systematic risk β_i	direct	+	+	$+^n$	+
	relative	+/−	+/−	+/−	+/−
size w_i	direct	+/(−)	+	+	+/−
	relative	$+^n$	$+^n$	$+^n$	$+^n$
Panel B: Sign of partial derivative in base case parameterization					
idiosyncratic risk σ_i	direct	−	+	+	+
	relative	−	+	+	+
systematic risk β_i	direct	+	+	+	+
	relative	+	+	+	+
size w_i	direct	+	+	+	+
	relative	+	+	+	+
Panel C: Sign of partial derivative in parameterization "one risky and dominant bank"					
idiosyncratic risk σ_i	direct	−	+	+	+
	relative	−	+	+	+
systematic risk β_i	direct	+	+	+	+
	relative	+	+	+	+
size w_i	direct	+	+	+	−
	relative	+	+	+	+
Panel D: Cases with neg. derivatives and references to Appendix A (analytic results)					
idiosyncratic risk σ_i	direct	Base case, A.9	A.5	A.5	A.4
	relative	Base case, A.10	A.11	A.11	A.11
systematic risk β_i	direct	A.2	A.6	A.7	A.6
	relative	Figure 3, A.13	A.12	A.12	A.12
size w_i	direct	A.3	A.9	A.9	Figure 4, A.8
	relative	A.15	A.14	A.14	A.14

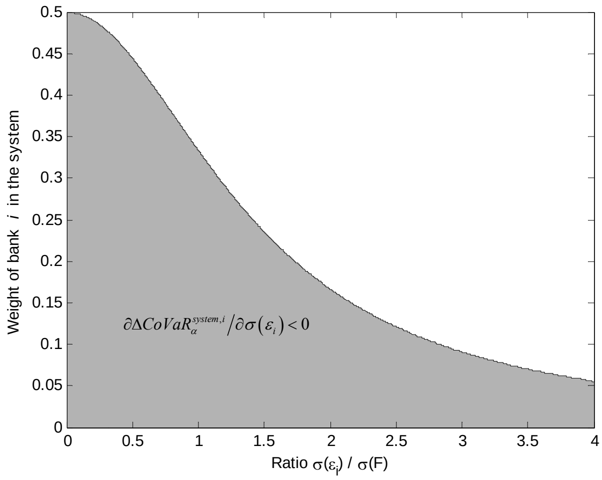

Figure 1: Regions in which ΔCoVaR is a falling function of idiosyncratic risk. We examine a system with $N = 50$ banks. We focus on the idiosyncratic risk of bank i, which has weight w_i in the system. Returns are described through

$$R_j = \beta_j F + \varepsilon_j, \quad R_S = \sum_{j=1}^{N} w_j R_j \text{ with } F \sim N\left(\mu, \sigma_F^2\right), \quad \varepsilon_j \sim N\left(0, \sigma_j^2\right).$$

All ε_j and F are independent. For the graph, banks are assumed to have a uniform β of 1. The horizontal axis is given by the ratio of idiosyncratic to systematic risk. The vertical axis plots w_i; other banks have a uniform weight $(1 - w_i)/49$. The figure shows those points as a gray area where the inequality $\partial \Delta CoVaR_\alpha^{system,i} / \partial \sigma_i < 0$ holds.

upper graph indicates, regardless of the bank's weight, that the ΔCoVaR always grows with β_i, as it should; in the appendix we show this to hold in general. However, the side effect (middle graph) can be so strong that the ΔCoVaR relative to other banks decreases. This is observed in the graph at the bottom in cases where bank i has a very large weight in the system. If a bank were increasing its exposure to systematic risk, a hypothetical ΔCoVaR-based systemic risk charge would therefore well "punish" this bank, but even more so its competitors.

To reveal the conditions under which the sensitivity to β_i is so strange, we perform an approximation, assuming that all banks have negligible weights, except bank i. Defining $\kappa \equiv w_i / (1 - w_i)$ as a monotonic function of w_i, $\beta_* \equiv (1 - w_i) \sum_{j \neq i} w_j \beta_j$ as the other banks' weighted average of systematic risks, and $\gamma \equiv \beta_i / \beta_*$, we obtain

$$\frac{\partial}{\partial \beta_i} \left[\frac{\Delta CoVaR_\alpha^{S,i}}{\Delta CoVaR_\alpha^{S,j}} \right] \propto ... \approx 1 - \kappa \left(\kappa \frac{\sigma_i^2}{\sigma_F^2 \beta_*^2} + \gamma (\kappa \gamma - 1) \right), \tag{11}$$

which is derived in Appendix A.13. This expression tends to be negative (i) when bank i is large relative to the others, (ii) when its idiosyncratic risk is large relative to the other banks' average systematic risk, given by $\sigma_F \beta_*$, and (iii), when its β_i is above the average. In other words, if the ΔCoVaR were used to define a systemic risk charge, the effect of β_i on a bank's own risk charge would be weaker than that on its competitors' charge, and this disproportion would be most pronounced if the bank was already riskier and larger than the others. In a particularly competitive environment, the relative effect might be beneficial enough to outweigh the increased costs involved with the own risk charge. In this way, ΔCoVaR could set a risk-increasing incentive.

A similar relative effect can be observed for the other measures. Using the same numerical example as for Figure 3, where σ_i is high and σ_F is low, the ratio[6] SRM_i / SRM_j is falling in β_i

[6] Here, SRM stands for exposure ΔCoVaR, MES or Beta.

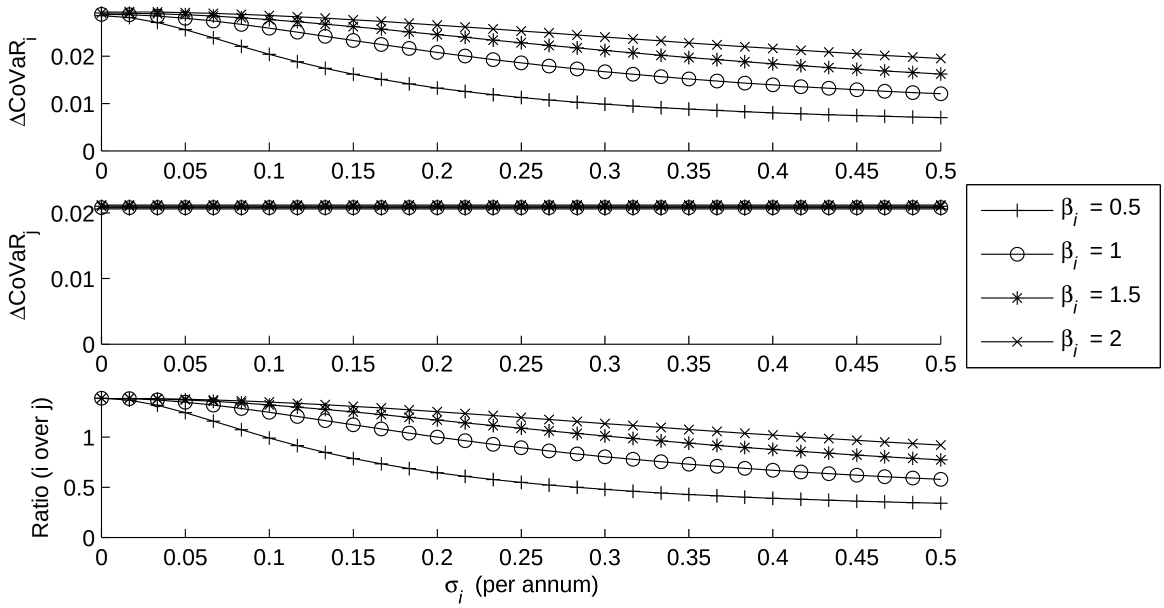

Figure 2: How ΔCoVaR responds to idiosyncratic risk. We examine the same system of $N = 50$ banks as in Figure 1, with the following additional assumptions. Banks have equal weights 1/50. Except bank i, all banks are uniform. They have a $\beta_j = 1$ and a standard deviation of idiosyncratic risk of 0.2 on an annual basis, the same as that of the systematic factor. On the x axis, the idiosyncratic risk σ_i of bank i varies from 0 to 0.5. Lines in the upper graph, each for a certain exposure β_i to the systematic factor, plot the ΔCoVaR of bank i on a daily basis at a quantile level of 0.01. The middle graph shows the ΔCoVaR for one of the other banks. The bottom graph shows the ΔCoVaR for bank i divided by that of another bank.

---

if the weight of bank i is exceptionally high. The partial derivative calculated in Appendix A.12 gives the same result: assuming for simplicity that the weight of bank j in the system can be neglected, we find

$$\lim_{w_j \to 0} \frac{\partial}{\partial \beta_i} \left[\frac{SRM_i}{SRM_j} \right] \propto 1 - \frac{w_i^2 \sigma_i^2}{\bar{\beta}^2 \sigma_F^2} .$$

Hence, β_i will have a negative relative effect under the condition that idiosyncratic risk, multiplied by the weight of bank i, exceeds average systematic risk, given by $\bar{\beta}\sigma_F$. There is a commonality with ΔCoVaR in that the relative effect on all four systemic risk measures tends to be negative if size and idiosyncratic risk are high while average systematic risk is low.

While the effect of a bank's size on ΔCoVaR is worth a look as well, we discuss it without technical details; these are found in Appendix A.3. Recall that we treat size independently of its particular meaning in this part of the paper, be it total assets or market capitalization. Assuming the size of bank i to change while those of the others remain unchanged, we observe an ambiguous effect in the following sense. On the one hand, the ΔCoVaR would increase with size under fairly general conditions, e.g., if β_i is not smaller than the average β of the other banks. If, on the other hand, the bank has sufficiently low systematic and idiosyncratic risk, the sensitivity can switch so that the bank would be assigned less systemic risk, and this would hold for any size. We can show, however, that the whole system would necessarily become less volatile then. The ΔCoVaR thus sets a correct incentive insofar as growth would be rewarded only if this also decreases system volatility.

We now discuss a further ambiguous case, which is the direct effect of size on beta. The general formula for the partial derivative is too complicated to identify simple conditions under which it is positive. We therefore make the simplifying assumption that only the bank of interest

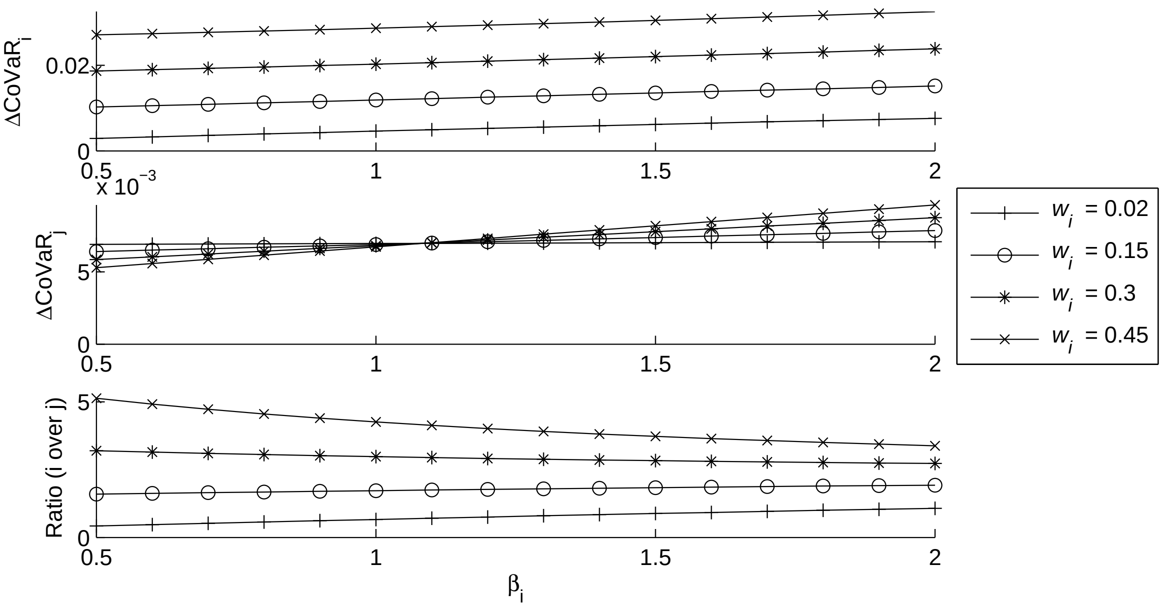

Figure 3: How ΔCoVaR responds to systematic risk. We examine the same system of 50 banks as in Figure 2, with the following modifications: $\sigma_i = 0.4$; $\sigma_F = 0.1$ (both p.a.); $\beta_j = 1$ for $j \neq i$. On the x axis, the systematic risk β_i of bank i varies from 0.5 to 2. The lines in the upper graph, each for a certain exposure weight w_i in the system, plot the ΔCoVaR of bank i on a daily basis at a quantile level of 0.01. The middle graph shows the ΔCoVaR for one of the other (uniform) banks. The bottom graph shows the ΔCoVaR for bank i divided by that of another bank.

i has a non-negligible weight in the system, whereas all other banks are infinitesimally small. Then they do not contribute idiosyncratic risks to the system anymore. As the resulting formula is still difficult to interpret, we consider the limiting case where also w_i is small. Formula (A-8) in Appendix A.8 gives

$$\lim_{w_i \to 0} \frac{\partial\, beta_i}{\partial w_i} \propto \sigma_i^2 - \beta_i \left(\beta_i - \beta_*\right) \sigma_F^2 \,,$$

where β_* is the average exposure of the other banks to systematic risk. The derivative is negative if the bank's exposure to systematic risk is above the average and the idiosyncratic risk is comparably small. Even for arbitrarily large σ_i there is always a β_i above which the derivative becomes negative. If a systemic risk charge were based on $beta_i$ without further corrections, the direct effect of size would be particularly undesirable since exactly those banks would be rewarded for growth that bear more systematic risk than the others, while this growth would increase the variance of the system return; the latter is shown in Equation (A-9) in Appendix A.8.

Numerical examples suggest similar effects in the general case where w_i is positive and the idiosyncratic risks of other banks are non-negligible. Figure 4 illustrates the extent to which a bank might benefit from growth if beta were used as a systemic risk charge. We assume the base case parameters for the factor as well as for idiosyncratic risk, and take all banks except bank i to have a β of 1. Bank i's weight in the system varies from 0 to 0.5 so that it can become very dominant relative to the other 49 banks with weight $w_j = (1 - w_i)/49$. For different exposures β_i to systematic risk, the figure shows how $beta_i$ depends on bank $i's$ weight. Up to $\beta_i = 1.5$, the bank would not benefit from growth – in sharp contrast to the case $\beta_i = 3$.[7] If bank i grows from

[7]This parameter, corresponding to betas between 2 and 3, appears fairly large. Here are some empirical observations for comparison. Stiroh (2006) reports an average beta for banks of 0.45 at a standard deviation of 0.42. He finds a maximum beta of 3.41. Note that these figures are betas for an index which covers all sectors,

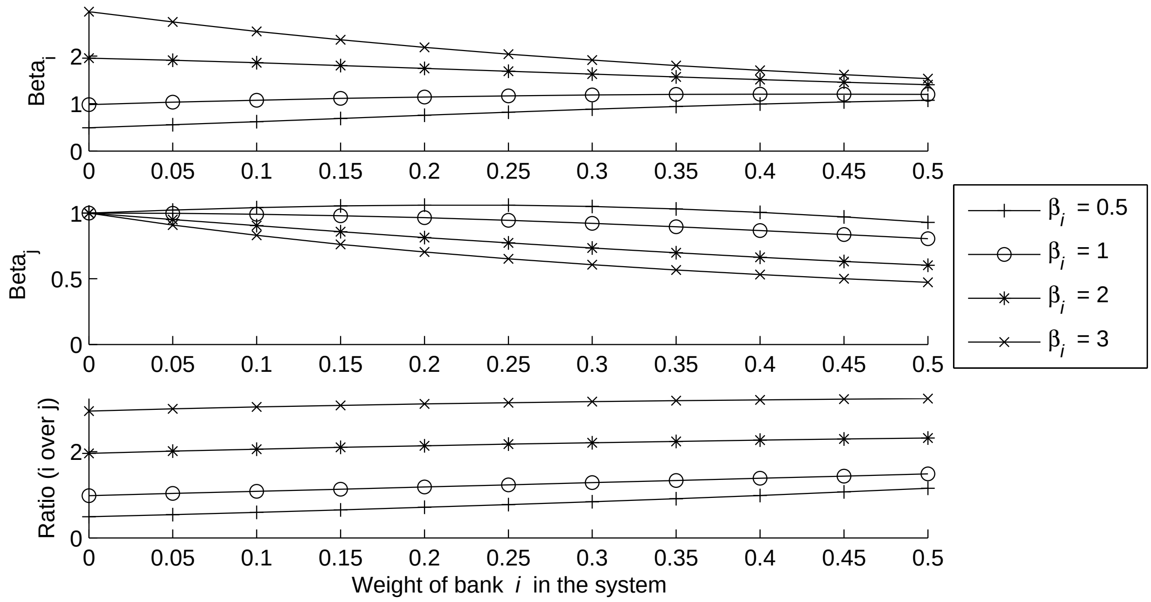

Figure 4: How beta responds to size. We examine the same system of $N = 50$ banks as in Figure 2. On the x axis, the weight of bank i varies from 0 to one half. The lines in the upper graph, each for a certain exposure β_i to the systematic factor, plot the beta of bank i, which is the regression coefficient b_i in $R_i = a_i + b_i R_S + u_i$. It is calculated without estimation error according to equation (9). The middle graph shows the beta for one of the other (uniform) banks. The bottom graph shows the beta for bank i divided by that of another bank.

zero to 29% of the system, ceteris paribus, it lowers its beta by a full integer. At the same time, the index becomes more volatile; $\sigma(R_S)$ is increased from 20.2 to 32.2 percent (p.a.). Turning to the side effect (middle graph), we observe that the betas of the other banks decrease as well.[8] This decrease is so strong that the lines for the ratios $beta_i/beta_j$ (bottom graph) are nearly flat. However, in Appendix A.14 we analytically show that the relative size effect on beta is still positive.

Observing that the relative size effect on beta is appropriate, while the direct effect is not always so, we could try to "repair" the direct effect by introducing a sensible scaling factor. Ad hoc, the first step might be a normalization, which is obsolete, however, since the weighted average of betas equals 1 by construction. Second, we might scale the beta up by a measure of system risk, e.g. $\sigma(R_S)$. Doing so is very much in line with the "component value-at-risk" approach introduced by Gauthier et al. (2012) as one of their candidate rules for capital allocation.

Looking up (8) and (9) for representations of the other systemic risk measures, we find our scaled beta to be $\text{cov}(R_S, R_i)/\sigma(R_S)$, which is essentially the same as exposure ΔCoVaR and MES, apart from constant factors and a typically small offset in the case of the MES. As they both get their sensitivities right, we conclude that rescaling the beta would indeed help but end in a concept equivalent to the exposure ΔCoVaR.

To summarize, a fairly clear message can be taken away from our analysis of sensitivities. If a systemic risk charge for banks were defined as a monotonic function of ΔCoVaR, the charge

unlike in our setup where only banks are included. A regression on an index of banks only would give an average beta of 1, so that we might expect standard deviations for the betas at the order of $0.42/0.45 = 0.93$, as a rough approximation. Betas above 2 or even 3 thus do not appear implausible. Related year-specific beta estimates reported by Baele, De Jonghe, and Vander Vennet (2007) have a mean of 0.61 and annual standard deviations around 0.50, which would suggest a standard deviation of betas around 0.82 ($= 0.5/0.61$) for a bank based index.

[8] The finding that *all* betas decrease may be surprising at first glance since their weighted average must equal 1. It can be explained by noting that the high-beta bank gains weight along with the fall in betas.

could set adverse incentives w.r.t. banks' risk parameters under many realistic conditions. MES and ΔCoVaR are nearly equivalent in the linear normal setup. All measures could set a wrong incentive w.r.t. the exposure to systematic risk if the bank is very dominant in the system. Beta can behave adversely w.r.t. size. If rescaled appropriately, beta boils down to a measure equivalent to the exposure ΔCoVaR.

3.3 Robustness to distributional assumptions

So far, we have assumed returns to follow a multivariate normal distribution. It yields tractable solutions but has limitations. In contrast to our previous assumptions, empirical returns may exhibit heavy tails as well as tail dependence, features to which systemic risk measures might react sensitively because they focus on tail behavior. Heavy tails can be due to time-varying volatilities, another empirical feature of returns that we have not modeled so far. Furthermore, the choice of asset vs. equity returns as the basis for systemic risk measures is something that we cannot consistently study with a multivariate normal framework. In a standard structural model of debt, for example, asset values follow a lognormal distribution, while equity – behaving like a call option on assets – has a non-lognormal return distribution.

To study whether the previous results are robust to variations in distributional assumptions, we first use a structural model to generate theoretically consistent asset and equity returns. Equity returns in this model exhibit time-varying volatilities and sensitivities, and are therefore heavy-tailed and tail dependent. We then complement this robustness check by replacing the multivariate normal distribution of the previous sections with a multivariate t, through which we can easily vary the degree of tailedness and tail dependence.

For the first robustness test we extend a structural model developed by Collin-Dufresne and Goldstein (2001) to multidimensional processes. We prefer this model as it is one of the few that generate stationary returns both for assets and equity. While Appendix C describes our modeling approach in detail, the main characteristics are as follows:

Asset values follow geometric Brownian motions. They are correlated through exposures to a common factor, which is also a geometric Brownian motion. Drifts are derived from a market model for asset values where, as in the CAPM for equity, risk premia for individual asset returns are proportional to the coefficient β_i, which links the asset return of bank i with the common factor.

While the lognormal asset returns have static parameters, debt is continuously adjusted by bank management to keep the equity ratio (defined as equity value over asset value) close to a strategic target level. The strength of the adjustment is governed by a speed-of-adjustment parameter and the current difference between the equity ratio and its target. In consequence, the equity ratio is a mean reverting process. As the ratio fluctuates around the target level, the instantaneous volatility of equity is driven up and down since its approximate relation to the (constant) asset volatility is $\sigma\left(R_{t,\text{equity}}\right) \approx \sigma\left(R_{t,\text{assets}}\right) \times V_t/E_t$, where V_t/E_t is the reciprocal equity ratio. This stochastic volatility makes the model similar to the empirical approach taken by Brownlees and Engle (2012) and Acharya et al. (2012) to analyze systemic risk measures, but there are also essential differences.[9]

We use simulations[10] to perform three tests that closely correspond to the scenarios in Fig-

[9]While their model is richer in that it includes dynamic correlations between systematic and idiosyncratic shocks to equity returns, asset returns are not explicitly modeled. By contrast, we put weight on the consistency of asset and equity returns, which are both stationary. Default events are explicitly modeled.

[10]We run 10 million independent simulations of bank and system returns. In CoVaR calculations, simulated returns would not fall into the conditioning event because its probability is zero. Instead we condition on a superset of positive probability. The $COVAR_{0.01}^{S|R_1=VaR_{0.01}^1}$, for example, is determined as follows: select the simulation runs in which the return R_1 lies between the 0.8% and 1.2% quantile of R_1, and determine the 1%

Table 2: Effects of risk parameters on systemic risk measures in a non-normal setting.
We use simulations to analyze sensitivities of systemic risk measures to certain risk parameters in a multivariate extension of the structural model of Collin-Dufresne and Goldstein (2001). Details are described in Appendix C.
Lognormal asset returns are drawn from geometric Brownian motions that are correlated through a common factor. Equity returns are derived from a stochastic differential equation, assuming debt is continuously adapted towards a target leverage ratio. Both returns are stationary, conditional on banks' survival.
A *direct effect* is understood as the change in a bank's own systemic risk measure owing to a change in one of the bank's parameters. The *relative effect* describes changes in the ratio between the systemic risk measures of two banks, e.g. MES_i/MES_j, while the risk parameter of bank i changes.
A plus (minus) sign indicates that the systemic risk measure / the ratio of two measures grows (falls) with the parameter. If both signs appear, the sign can be positive or negative, depending on other parameters. Gray cells +/− mark cases where the outcome is different from its counterpart in Panel A of Table 1 (normal linear setup).
The base case parameters (p.a. for drift, volatility and mean reversion) are set to $N = 50$, $\mathbf{E}(F) = 0.03$, $\sigma_F = 0.05$, $\sigma_j = 0.04$, $\beta_j = 1$, $w_j = 1/50$, target log debt ratio $\bar{l}_j = -0.1$, mean reversion $\lambda_j = 2.38$, for all j; quantile level $\alpha = 0.01$ for the CoVaR measures and 0.05 for the MES.
For the effect of idiosyncratic risk, σ_i varies from 0.01 to 0.1. Monotonicity in σ_i is checked for β_i between 0.5 and 2.
For the effect of systematic risk, β_i varies from 0.5 to 2. Monotonicity in β_i is checked for a weight w_i between 0.02 (equal share) and 0.45. In this exercise we set $\sigma_i = 0.08$ and $\sigma_F = 0.025$ to provoke the effect found in the normal model. For the effect of a bank's weight in the system, w_i varies from 0.02 (equal share) to 0.45. The effect of w_i is checked for β_i between 0.5 and 3.

Parameter	Effect	Return type	ΔCoVaR	Exp. ΔCoVaR	MES	Beta
idiosyncratic risk σ_i	direct	assets	−	+	+	+
		equity	−	+	+/−	+
	relative	assets	−	+	+	+
		equity	−	+	+/−	+
systematic risk β_i	direct	assets	+	+	+	+
		equity	+	+	+	+
	relative	assets	+/−	+/−	+/−	+/−
		equity	+/−	+/−	+/−	+/−
size w_i	direct	assets	+	+	+	+/−
		equity	+	+/−	+	+/−
	relative	assets	+	+	+	+
		equity	+/−	+/−	+	+

ures 2–4. The risk parameters for asset returns are broadly consistent with Moody's KMV asset volatilities of banks and the values found by Memmel and Raupach (2010); from there we also take values for mean reversion speed and target leverage. Table 2 presents parameters and results.

In the first test, representing the counterpart to the analysis illustrated in Figure 2, we vary the idiosyncratic risk of asset returns and check the direct and relative effect on the systemic risk measures. This is done for different β_i. Both for asset and equity returns we observe that an

quantile of R_S for this selection. The MES is determined as the average simulated return of a bank under the condition that the simulated system return is below its 5% quantile. The beta is estimated through a regression of a bank's simulated returns on the simulated system returns.

increase in idiosyncratic risk lowers ΔCoVaR, which is the same result as in the normal model. Importantly, bank i's equity returns get heavier distribution tails when idiosyncratic risk rises[11], meaning that the decrease in ΔCoVaR goes along with both increased individual variance and heavier tails. Here we find negative sensitivities where they did not exist in the normal model: for high β_i, there is both a direct and relative negative effect on MES. Interestingly, no such effect is observed for asset returns, which suggests that tail thickness, as the outstanding difference in the return distributions, plays a role here.[12]

Second, similar to Figure 3, we test the sensitivity to β_i, using different weights of bank i in the system. To examine whether we obtain similar effects as in the normal setup, other risk factors are given values derived from (11) to provoke negative sensitivities.[13] Both for asset and equity returns, an increasing β_i has a negative relative effect on all four systemic risk measures if the weight of a bank is very high (0.45). This finding conforms to the results of the normal model. Again, the increase in β_i is accompanied by a (moderate) increase in tail thickness.[14]

In a third test, similar to Figure 4, we vary the weight w_i of bank i in the system. Using different values for β_i, the results of the normal model are confirmed insofar as an increase of the weight can lower the beta if the factor sensitivity β_i is high, which holds for asset and equity returns. Further negative sensitivities appear that were nonexistent in the normal model: for high β_i, there is a negative relative size effect on ΔCoVaR and both a direct and relative negative effect on exposure ΔCoVaR. As in the first test, no such effect is observed for asset returns, which seems to confirm that tail thickness matters.

What we observe when w_i grows while $\beta_i = 2$ is that a highly leptokurtic return (with a constant kurtosis of 8.5) increasingly shapes the index return, the kurtosis of which grows from 3.4 to 5.9. Hence, the riskiness of the index does not only rise for its increased volatility but also for a heavier loss tail. Both effects go along with a negative relative effect on beta and either ΔCoVaR version.

We conclude this robustness test with the observation that all undesired sensitivities found in the normal model are confirmed for lognormal asset returns and for equity returns with heavier tails. However, further problematic sensitivities show up if distribution tails are thicker.

As an alternative way of generating heavy-tailed returns, we consider a multivariate t distribution. To draw returns from a multivariate t with ν degrees of freedom, we first simulate multivariate normal returns according to (2) and then scale these returns with $((\nu - 2)/w)^{0.5}$, where w is a random draw from a chi-square distribution with ν degrees of freedom.[15] The random scaling factor is the same for the N banks of one trial, but differs across trials.

As before, we determine the systemic risk measures using simulations.[16] We consider the degrees of freedom ν to be 4, 5, 6, 8, or 10 for different parameterizations of (2) and thus obtain

[11]When σ_i varies from 0.01 to 0.1, the kurtosis of the equity return grows from 3.10 to 6.66 (assuming $\beta_i = 0.5$) or from 5.47 to 8.38 (assuming $\beta_i = 2$). The equity index return has a fairly stable kurtosis around 3.4.

[12]To exclude that the difference in the effect on the systemic risk measures is simply due to differences in the general volatility level, we do the following exercise: equity returns are rescaled after simulation such that they have the same daily standard deviation as their corresponding asset returns. After rescaling, the exposure ΔCoVaR still exhibits the same negative sensitivity to size (which does not exist for asset returns). As the correlation matrices of equity and asset returns are also very similar, tail thickness is the only plausible remaining explanation for the fact that the negative effect is observed with equity returns only.

[13]The equations (11) and (A-10) do not actually apply here but may give an indication. In the test, σ_i is doubled from 0.04 to 0.08 (p.a.), while σ_F is halved from 0.05 to 0.025.

[14]When β_i varies from 0.5 to 2 (assuming the highest weight for bank i), the kurtosis of the equity return grows from 5.54 to 6.15. The equity index return exhibits kurtosis values between 4.82 and 5.08.

[15]The t distribution can be characterized as a normal variance mixture with mixing factor $(\nu/w)^{0.5}$. We then divide by the standard deviation of the t distribution $(\nu/(\nu - 2))^{0.5}$ to obtain the same standard deviation as in the normal model.

[16]For the t distributed returns we perform 100 million antithetic simulations. Systemic risk measures are calculated as before.

a large number of results, which we do not report in full here. They are similar to the ones from the structural model in that key results from the previous section continue to be found in typical cases, and that there are some variations. For example, when we start with the base case parameters and then increase idiosyncratic risk from 0.2 to 0.3, systematic risk from 1 to 1.5 or the weight of one bank from 0.02 to 0.25, the simulations show negative effects on the risk measures in each of the two cases where Panel B of Table 1 has a negative sign; in the cases where the effect of Table 1 has a positive sign, the simulated effects are also positive except for $\nu = 4$, which shows additional negative effects for ΔCoVaR and exposure ΔCoVaR. There are other situations in which the degrees of freedom have an influence on the sign of effects. When we move from the base case parameterization to a lower idiosyncratic risk of 0.1, ΔCoVaR increases for $\nu \geq 5$, consistent with the negative sensitivity found in the normal model, while it decreases for $\nu = 4$. This observation does not question the general conclusion, though. As just described, lowering the degrees of freedom to small values does not only make undesired sensitivities disappear but also makes new ones appear.

Summing up, replacing the multivariate normal with other distributions leads to results that are more complex than before, but the key conclusion of the previous section is confirmed: there are non-exotic situations in which a parameter change that increases the risk of the system lowers the estimated systemic risk contribution of the bank for which the parameter was changed.

4 Systemic risk measures in the contagion case

After studying linear return relationships within a simple one-factor model, we now turn to examining contagion effects. An overview of different contagion definitions is given in Pericoli and Sbracia (2003). The main definition that we examine is one in which contagion is brought about by spillovers of idiosyncratic shocks.

Assume that the returns of the banks and the system evolve according to

$$R_i = \beta_i F + \varepsilon_i + \sum_{j \neq i} \gamma_j I_{\{\varepsilon_j < \kappa\}} \varepsilon_j, \quad R_S = \sum_i w_i R_i \,. \tag{12}$$

That is, there can be contagion from one bank to other banks in the system. If bank j is afflicted by a realization of idiosyncratic risk that is worse than κ, other banks are partially affected, too. As in the previous section, we assume that F and all ε_j are independent normal variates. Since the dependence structure is now considerably more involved, we resort to Monte Carlo simulation to derive statements about systemic risk measures. Based on 100 million simulated returns, the measures are estimated as described in Footnote 10.

As before, the analysis is conducted using assumptions typical of daily returns. In the base case, banks are equally weighted ($w_i = 1/N$), factor betas are uniformly set to one, and the following per-annum drift and volatility and parameters are chosen: $\mathbf{E}(F) = 0.05$, $\sigma_F = 0.2$, and $\sigma_i = 0.2$ for all i. As before, dividing these parameters by 260 or $\sqrt{260}$ translates them to daily returns. The number of banks is set to $N = 50$.

For the first analysis, we assume that only bank 1 is infectious, i.e. $\gamma_j = 0$ for $j > 1$. We set the contagion threshold κ to -0.0204, which corresponds to the 5% quantile of ε_1. The only parameter that is left to specify is the contagion intensity γ_1. Figure 5 shows the risk figures that result if we vary γ_1 from 0 to 1 and keep the other parameters at the values just described. The way in which the risk measures depend on the contagion intensity differs markedly. Below we briefly describe and elucidate the observed patterns for each measure.

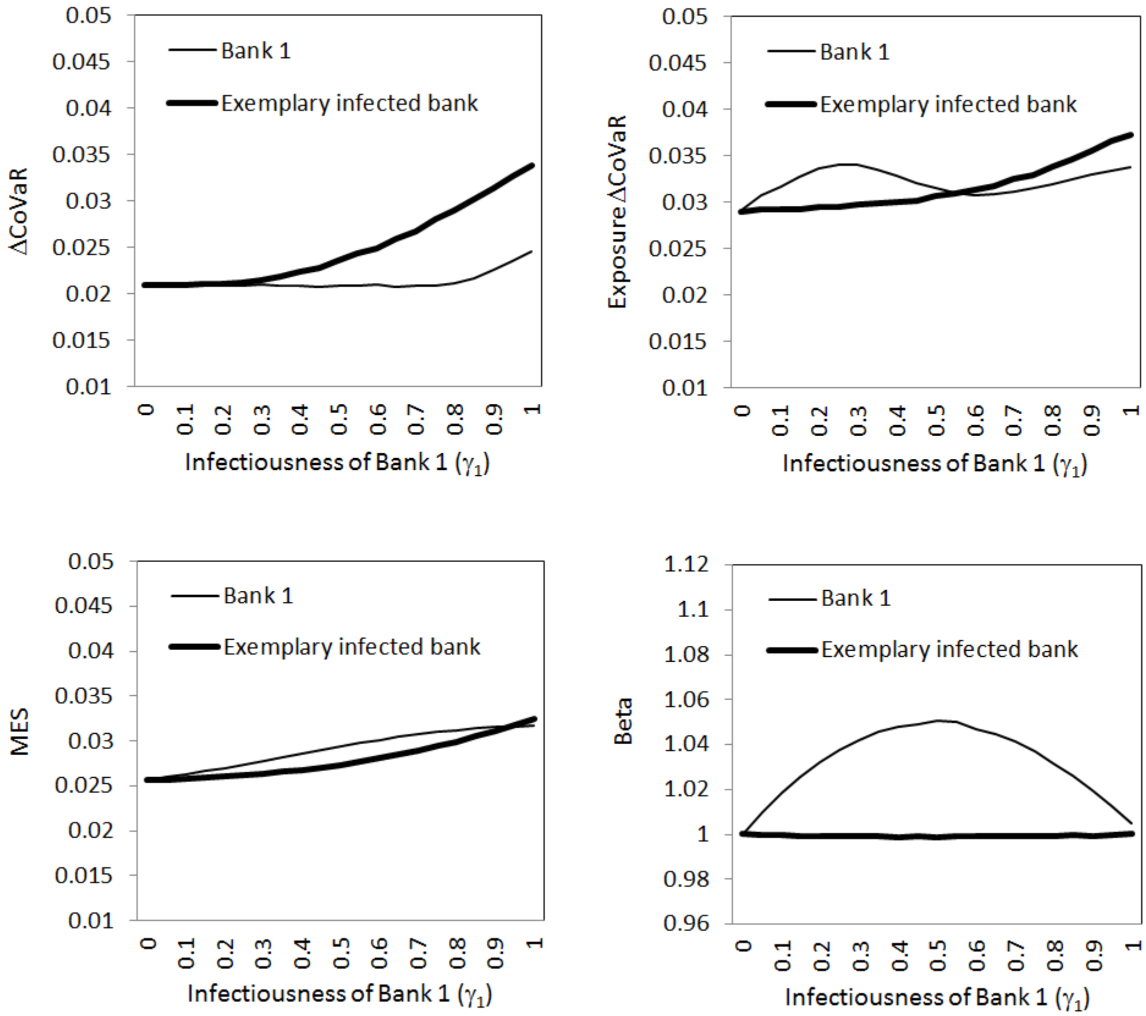

Figure 5: Risk measures in the contagion case. We simulate returns for N equally-sized banks. The banking system return is the average of bank returns. Bank returns are driven by a common factor F, idiosyncratic risk, and a spillover of idiosyncratic risk from bank 1 to the other banks:

$$R_1 = \beta_1 F + \varepsilon_1, \quad R_j = \beta_j F + \varepsilon_j + \gamma_1 I_{\{\varepsilon_1 < \kappa\}} \varepsilon_1 \text{ for } j > 1, \quad R_S = N^{-1} \sum_{j=1}^{N} R_j.$$

Parameters (p.a. for drift and volatility) are set to $N = 50$, $\kappa = -0.0204$, $\mathbf{E}(F) = 0.05$ (p.a.), $\sigma_F = 0.2$ (p.a.), $\sigma_i = 0.2$ (p.a.), $\beta_i = 1$ for all i. The infectiousness of bank 1 is varied from $\gamma_1 = 0$ to $\gamma_1 = 1$. The measures are estimated through Monte Carlo simulation with 100 million trials. CoVaR$_\alpha$ measures are computed on a daily basis for $\alpha = 1\%$, using observations between the $(\alpha - 0.2\%)$ and $(\alpha + 0.2\%)$ quantiles of the conditioning variable. The MES relates to $\alpha = 5\%$.

4.1 ΔCoVaR

To understand how ΔCoVaR is affected by the contagion intensity γ_1, we express the system return as a function of the return of an individual bank. For ease of exposition, we will incorporate the choice of uniform unit betas and equally-weighted banks that we made for the simulation. When applied to bank 1, the left-hand part of (12) then implies $F = R_1 - \varepsilon_1$, which can be plugged into the representation of R_S to eliminate the factor return F:

$$R_S = R_1 - \frac{N-1}{N}\left[1 - I_{\{\varepsilon_1<\kappa\}}\gamma_1\right]\varepsilon_1 + \frac{1}{N}\sum_{j=2}^{N}\varepsilon_j\,. \tag{13}$$

When we use (13) to study the system return conditional on a quantile of R_1, it is not enough to replace R_1 by its quantile. The conditional distribution of ε_1 also differs from the unconditional one. However, it is not affected by the contagion intensity γ_1, which facilitates the analysis.

Changing γ_1 can influence the ΔCoVaR through effects on both the 50% quantile and the 1% quantile of R_S. For the 50% quantile, effects will be relatively small because the probability that contagion occurs if R_1 is at its median is very small. This probability, which again is independent of γ_1, can be determined by exploiting the fact that the conditional distribution of ε_1 is normal (see Appendix B). For the parameter combination used here, contagion occurs with a probability of 1.00% if R_1 is at its median, and with a probability of 50.02% if R_1 is at its 1% quantile (see Equation (B-3) in Appendix B). For the purpose of understanding the patterns of Figure 5, it is therefore sufficient to focus on the 1% CoVaR.

Figure 5 shows that changes in the contagion intensity do not affect the ΔCoVaR until γ_1 reaches a value of around 0.75. This may seem surprising, given that contagion happens with a probability of over 50% once R_1 is at its 1% quantile. However, it does not necessarily follow that contagion events are crucial for the CoVaR, which is the 1% quantile of R_S conditional on R_1 taking some value. Equation (13) shows that one way of arriving at a low conditional realization of R_S is to have a very positive realization of ε_1. If ε_1 is positive, however, there is no contagion.

The extent to which low realizations of R_S are associated with contagion is illustrated in Figure 6. It plots the simulated R_S against the simulated ε_1 conditional on R_1 being near its 1% quantile.[17] With $\gamma_1 = 0.6$, none of the simulated cases of contagion is associated with a system return that is below its 1% quantile. In consequence, a change in the contagion intensity does not affect the ΔCoVaR. Moving on to $\gamma_1 = 0.8$ and $\gamma_1 = 1$, contagion increasingly matters for the 1% quantile of the system return. In consequence, a higher γ_1 leads to a higher ΔCoVaR because it magnifies the negative effects of contagion.

For an infected bank – here we take it to be bank 2 – we can derive:

$$R_S = R_2 - \frac{N-1}{N}\varepsilon_2 + \frac{1}{N}\left(1 - I_{\{\varepsilon_1<\kappa\}}\gamma_1\right)\varepsilon_1 + \frac{1}{N}\sum_{j=3}^{N}\varepsilon_j\,. \tag{14}$$

The direct effects of γ_1 and ε_1 that we discussed above now play a smaller role because they enter the equation with a factor of $1/N$ rather than $(N-1)/N$. However, there is an additional effect because the quantiles of R_2 also depend on γ_1. Increasing γ_1 lowers both the median and the 1% quantile of R_2, with the effect on the latter being more pronounced. This is the key factor behind the pattern shown in Figure 5: contagion increases the risk of the infected bank as well as the entire system.

Comparing the infectious and the infected bank, Figure 5 shows that contagion drives a wedge between the ΔCoVaR of the two banks, which increases with the strength of the spillovers. In the presence of contagion, ΔCoVaR assigns a larger systemic risk to the infected bank.

While this pattern is consistent with the discussion from above, we would like to provide another way of understanding it. In Figure 7, we visualize the conditional distribution of the system return by plotting it against the return of the infectious bank and the return of an infected

[17] Recall that selecting observations in which the return R_1 lies between its 0.8% and 1.2% quantile is our numerical procedure for conditioning on the 1% quantile of R_1.

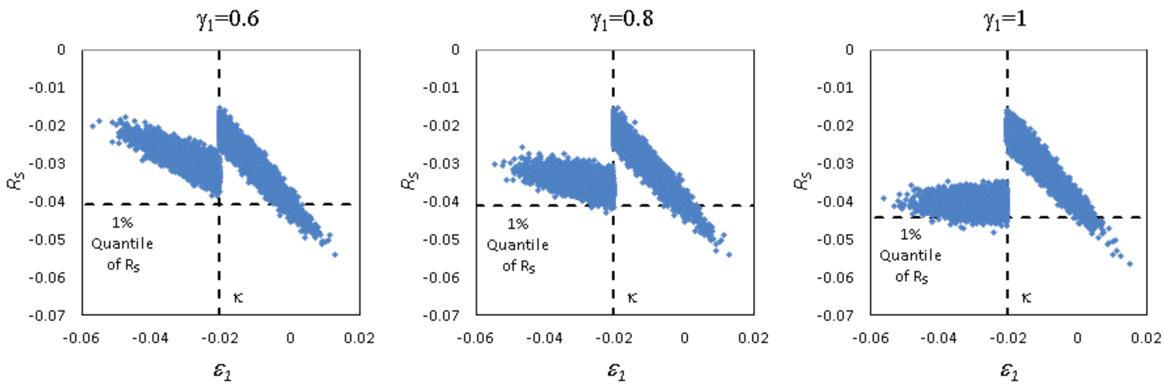

Figure 6: System return and idiosyncratic risk of the infectious bank if the return of the latter is at its 1% quantile. We simulate returns for the same system as in Figure 5 in which bad outcomes of ε_1 can spill over to returns of the other banks. Eliminating the systematic factor implies:

$$R_S = R_1 + \frac{N-1}{N}\left[I_{\{\varepsilon_1<\kappa\}}\gamma_1 - 1\right]\varepsilon_1 + \frac{1}{N}\sum_{j=2}^{N}\varepsilon_j.$$

In the chart, we plot simulated R_S against simulated ε_i for a selection of observations in which R_1 is in a close interval ($\pm0.2\%$) around its 1% quantile. The infectiousness parameter γ_1 is set to either 0.6, 0.8, or 1.

bank, respectively. We choose the contagion intensity γ_1 to be 0.75 and we also show scatterplots for two subsamples defined according to whether there is contagion ($\varepsilon_i < \kappa$) or not. Conditional on the bank returns being at their 1% quantiles, the system return has a lower mean in the case of the infectious bank.

While this tends to make the ΔCoVaR of the infected bank more extreme, there is a stronger effect working in the opposite direction. For the infectious bank, the conditional variance of the system return is relatively low. In the contagion case, the system return is highly correlated with the infectious bank because that bank's idiosyncratic risk has spread through the system. If there is no contagion, the conditional volatility is relatively low because it is then relatively likely that the system return has been brought about by a low factor return.[18]

The low conditional variance of the system return means that – according to ΔCoVaR – the systems appears to have a relatively low risk, conditional on the infectious bank being at its 1% quantile.

4.2 Exposure ΔCoVaR

Comparing the four measures in Figure 5, the exposure ΔCoVaR of the infectious bank exhibits the most complex pattern. At first, it becomes larger. Then it shrinks, but this new tendency is again reversed. Rearranging (13), we can examine how the return of the infectious bank depends on the system return:

$$R_1 = R_S + \frac{N-1}{N}\left(1 - I_{\{\varepsilon_1<\kappa\}}\gamma_1\right)\varepsilon_1 - \frac{1}{N}\sum_{j=2}^{N}\varepsilon_j\,. \tag{15}$$

Using this equation is now considerably more involved than above. A change in γ_1 affects the quantiles of R_S as well as the conditional distribution of ε_1. An inspection of the simulated

[18]If there is no contagion, the system return can be low because of a low factor realization or a low average realization of the infected banks' idiosyncratic risk. With 49 banks, however, the variance of the average idiosyncratic shock is very low, making it less likely to be the reason for a low system return.

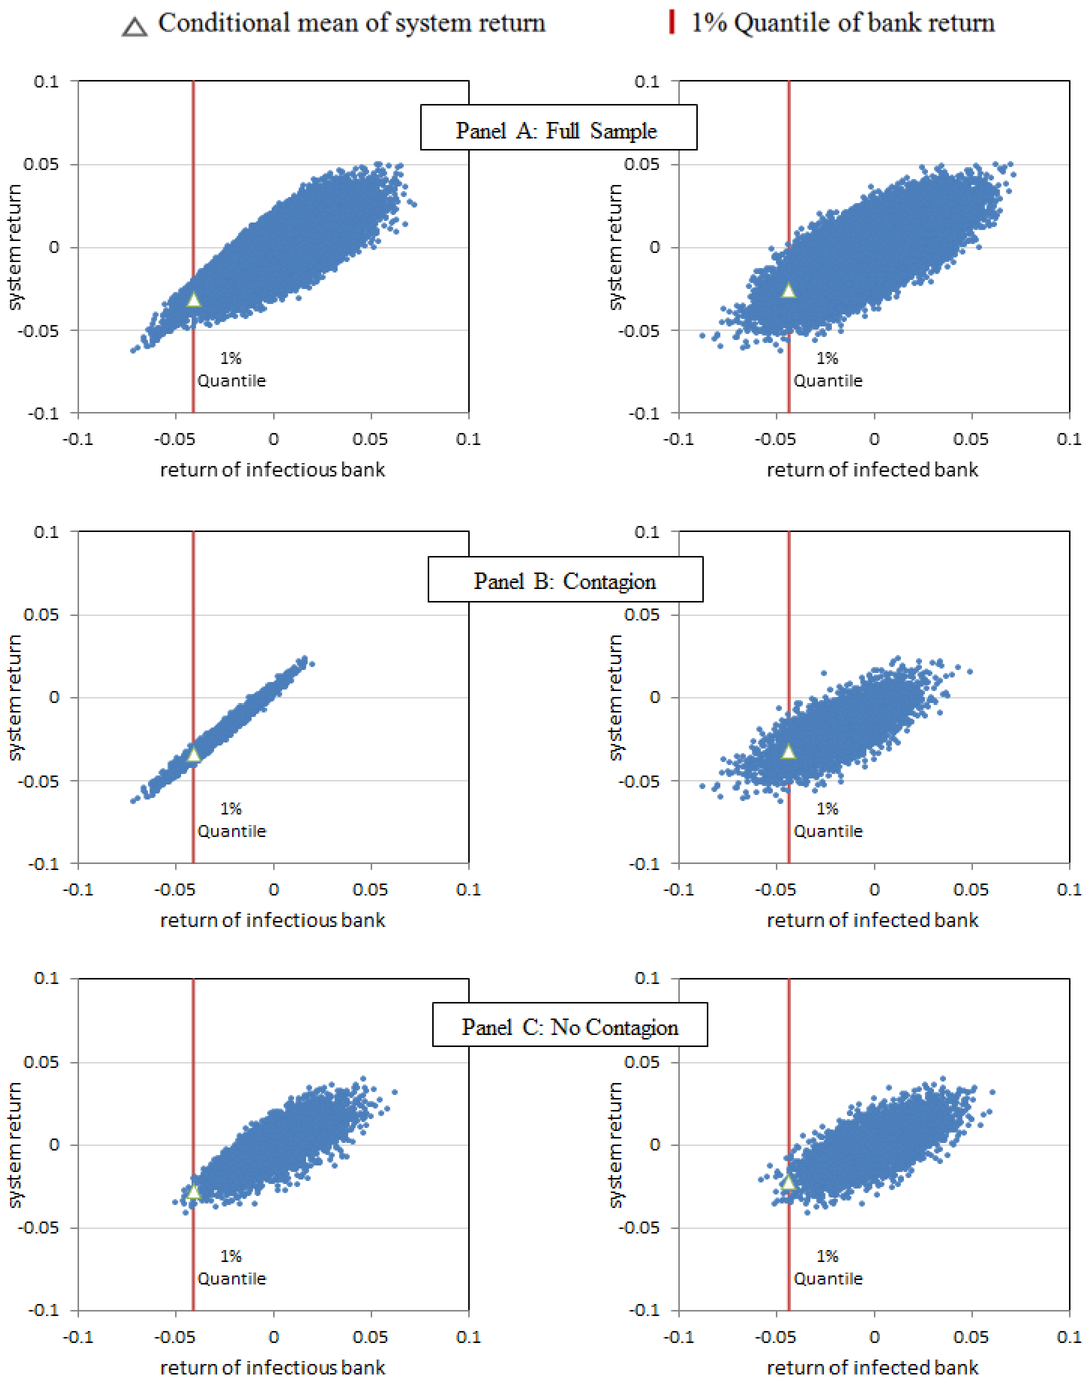

Figure 7: Simulated system returns versus returns of infectious and infected banks. We simulate returns for the same system as in Figure 5. The infectiousness of bank 1 is set to $\gamma_1 = 0.75$. We plot simulated system returns R_S against individual returns R_1 (infectious bank, on the left) and R_j (an infected bank, on the right). Panel A plots the full sample. Panel B contains only cases of contagion, where $\varepsilon_1 < \kappa$. Panel C contains cases of no contagion. The vertical red line marks the event that the individual bank return is at its 1% quantile, the event which $CoVaR_{1\%}$ conditions on.

CoVaR figures reveals which forces drive the observed patterns: the initial upward move goes back to a decrease in the $\text{CoVaR}_{50\%}$; this decrease is then overcompensated by changes in the $\text{CoVaR}_{1\%}$, which exhibits a trough-shaped behavior, with the bottom of the trough being close to $\gamma_1 = 0.5$. We explain those in turn.

An increase in γ_1 makes it less likely that R_S is at its median once contagion has occurred because contagion shifts R_S away from its median and this shift is stronger the larger γ_1.[19] In consequence, a higher γ_1 means that there will be fewer realizations with ε_1 below the contagion threshold if R_S is at its median. There is another effect having the same impact on $\text{CoVaR}_{50\%}$: an increasing γ_1 pulls $\left(1 - I_{\{\varepsilon_1<\kappa\}}\gamma_1\right)\varepsilon_1$ towards zero if there is contagion. Together, this makes the $\text{CoVaR}_{50\%}$ decrease, as it is defined as the negative of the 1% quantile of R_1 conditional on R_S being at its median. An opposite effect – an increase of γ_1 lowers the median of R_S – is relatively small.

What happens if we condition on the 1% quantile of R_S in order to determine the $\text{CoVaR}_{1\%}$? An increase in γ_1 now makes it *more* likely that it was contagion that led to the extreme realization of R_S. More contagion implies more extremely negative realizations of ε_1, but what matters for the $\text{CoVaR}_{1\%}$ is $\left(1 - I_{\{\varepsilon_1<\kappa\}}\gamma_1\right)\varepsilon_1$, which is pulled towards zero as γ_1 increases. To get an intuition why the second effect is stronger, consider the extreme case in which γ_1 equals 1. Then, the smallest value that $\left(1 - I_{\{\varepsilon_1<\kappa\}}\gamma_1\right)\varepsilon_1$ can take is κ, even though there will be many realizations with an ε_1 smaller than κ. The overall trough-shaped pattern in the exposure ΔCoVaR arises because this effect is at some point outsized by another effect of γ_1 which works in the opposite direction: increasing γ_1 lowers the 1% quantile of R_S which we condition on.

For the infected banks the pattern is less complex: the ΔCoVaR grows in γ_1. We can rearrange equation (14):

$$R_2 = R_S + \frac{N-1}{N}\varepsilon_2 + \frac{1}{N}\left(I_{\{\varepsilon_1<\kappa\}}\gamma_1 - 1\right)\varepsilon_1 - \frac{1}{N}\sum_{j=3}^{N}\varepsilon_j\,.$$

The effects of changes in the conditional distribution of $\left(1 - I_{\{\varepsilon_1<\kappa\}}\gamma_1\right)\varepsilon_1$ are now less important than in the case of the infectious bank studied above because the factor $1/N$ is much smaller than $(N-1)/N$. Changes in the conditional quantiles of R_2 are therefore driven by changes in the quantiles of R_S. Since the 1% quantile of R_S is more sensitive to changes in γ_1 than the 50% quantile, the exposure ΔCoVaR of bank 2 becomes larger.

4.3 MES

For the infectious bank, we again inspect:

$$R_1 = R_S + \frac{N-1}{N}\left(1 - I_{\{\varepsilon_1<\kappa\}}\gamma_1\right)\varepsilon_1 - \frac{1}{N}\sum_{j=2}^{N}\varepsilon_j\,.$$

MES and exposure ΔCoVaR are similar in that we condition on a tail event of the system return. Increasing γ_1 makes it more likely that contagion has occurred in the conditioning event that R_S is below its 5% quantile. A higher probability of contagion means that the conditional means of

[19] More precisely, contagion shifts the system return *down* so that some realizations of R_S will also be shifted (from above) towards the median. However, even with $\gamma_1 = 0$ the majority of "contagion" events has realizations of R_S below the median. If γ_1 rises, more returns with contagion are therefore pushed away from the median than are pushed towards it. This lowers the presence of contagion events among those having a return at the median, in total. There is also a slight compensating effect since the median of R_S is decreased, too; but the effect has lower order because the median is mainly driven by F and aggregate idiosyncratic risk, which both do not depend on γ_1.

both R_S and ε_1 are more negative. But the direct effect of γ_1 works in the opposite direction. Through $\left(1 - I_{\{\varepsilon_1<\kappa\}}\gamma_1\right)\varepsilon_1$, an increase in γ_1 makes the MES less extreme. The concave shape of the MES arises because the first effect dominates for small γ_1, while the second effect gains weight when γ_1 gets larger.

For an infected bank, the reasoning is very close to the one for the exposure ΔCoVaR:

$$R_2 = R_S + \frac{1}{N}\left(I_{\{\varepsilon_1<\kappa\}}\gamma_1 - 1\right)\varepsilon_1 - \frac{1}{N}\sum_{j=2}^{N}\varepsilon_j + \varepsilon_2 \,.$$

What matters most are changes in the quantile of R_S, which is driven down by increases in γ_1. The second term does not contribute much because it contains the factor $1/N$.

For γ_1 very close to 1, the MES indicates that the infected bank is riskier than the infectious bank. This is easiest understood if infection occurred and $\gamma_1 = 1$. In this case, the equations for R_1 and R_2 differ only in ε_2, which adds to R_2 by $+\varepsilon_2$ while it adds to R_1 by $-\varepsilon_2/N$. Conditional on $R_S < Q_S^\alpha$, the expected value of ε_2 is negative because it also contributes to the system return.

4.4 Beta

To understand why the beta of the infectious bank with respect to the system is hump-shaped in the contagion intensity γ_1, we study the covariance conditional on infection. (Without infection, we are in the standard one-factor case, irrespectively of the value of γ_1.) In (13) we replace R_1 by $F + \varepsilon_1$ to obtain

$$R_S = F + \left[\frac{N-1}{N} I_{\{\varepsilon_1<\kappa\}}\gamma_1 + \frac{1}{N}\right]\varepsilon_1 + \frac{1}{N}\sum_{j=2}^{N}\varepsilon_j \,.$$

Conditional on $\{\varepsilon_1 < \kappa\}$, the covariance is

$$\begin{aligned} \operatorname{cov}\left(R_1, R_S | \varepsilon_1 < \kappa\right) &= \operatorname{cov}\left(F + \varepsilon_1, F + \left[\frac{N-1}{N}\gamma_1 + \frac{1}{N}\right]\varepsilon_1 + \frac{1}{N}\sum_{j=2}^{N}\varepsilon_j \,\middle|\, \varepsilon_1 < \kappa\right) \\ &= \sigma_F^2 + \left[\frac{N-1}{N}\gamma_1 + \frac{1}{N}\right]\sigma_{1,\text{cont}}^2 \,, \end{aligned}$$

where $\sigma_{1,\text{cont}}^2$ is the $\{\varepsilon_1 < \kappa\}$-conditional variance of ε_1. Similarly, we obtain:

$$\sigma^2\left(R_S | \varepsilon_1 < \kappa\right) = \sigma_F^2 + \left[\frac{N-1}{N}\gamma_1 + \frac{1}{N}\right]^2 \sigma_{1,\text{cont}}^2 + \frac{1}{N^2}\sigma_j^2 \,.$$

The unconditional covariance, which matters for the beta, is also affected by the covariance of the conditional means. The mean of R_1 conditional on contagion is $\mathbf{E}(F) + \mathbf{E}(\varepsilon_1 | \varepsilon_1 < \kappa)$ and does not depend on γ_1, while the mean of R_S conditional on contagion is

$$\mathbf{E}(F) + \left(\frac{N-1}{N}\gamma_1 + \frac{1}{N}\right)\mathbf{E}(\varepsilon_1 | \varepsilon_1 < \kappa) \,,$$

which does depend on γ_1. In consequence, the covariance of the conditional means is linear in γ_1.

The observation that the covariance of R_S and R_1 is linear in γ_1 while the variance is quadratic in γ_1 explains the hump-shaped relationship between the beta of the infectious bank 1, which is $\operatorname{cov}(R_1, R_S)/\sigma^2(R_S)$, and the contagion intensity γ_1.

For the infected banks, and again conditional on contagion, we observe that the covariance is quadratic in γ_1:

$$\begin{aligned}\operatorname{cov}(R_2, R_S|\varepsilon_1 < \kappa) &= \operatorname{cov}\left(F + \varepsilon_2 + \gamma\varepsilon_1, F + \left[\frac{N-1}{N}\gamma_1 + \frac{1}{N}\right]\varepsilon_1 + \frac{1}{N}\sum_{j=2}^{N}\varepsilon_j \middle| \varepsilon_1 < \kappa\right) \\ &= \sigma_F^2 + \left[\frac{N-1}{N}\gamma_1^2 + \frac{1}{N}\gamma\right]\sigma_{1,\text{cont}}^2 + \frac{1}{N}\sigma_2^2 .\end{aligned}$$

The same holds for the covariance of the conditional means because the conditional mean of an infected bank's return as well as the mean of the system return depend on γ_1. In the definition of beta, both the numerator and the denominator are thus quadratic in γ_1, which explains the almost flat relationship between γ_1 and the beta of an infected bank.

4.5 Robustness

To examine whether the results are robust to the parameter choices and distributional assumptions made above, we examine several modifications:

[1] By making one bank infectious as we do in the base case, the infected banks have a lower mean return and a higher volatility than the infectious bank. To correct for this effect we add a constant to the infected banks' returns and change their idiosyncratic volatility such that they have the same expected return and overall volatility as the infectious bank. The size of the adjustments can be derived using standard results for truncated normal distributions.

[2] We change the contagion threshold from –0.0204 to –0.0289, which corresponds to the 1% quantile of the idiosyncratic risk component.

[3] We change the contagion threshold from –0.0204 to –0.0383, which corresponds to the 0.1% quantile of the idiosyncratic risk component.

[4] Infectious banks might be larger than infected banks. We therefore change the weight of the infectious bank from 2% in the base case to 25%. The other weights are changed to $(1 - 0.25)/49$.

[5] Infectious banks might have a larger exposure to the common factor. We therefore change the factor beta of the infectious bank from $\beta_1 = 1$ to 1.25.

[6] In the base case, there is just one infectious bank. In this variation, the number of infectious banks is increased to five, each of them having the same contagion intensity.

[7] Instead of being i.i.d. normal, we assume the factor return F to follow a GARCH(1,1) process with the same unconditional volatility as in the base case. We choose $\delta = 0.04$ and $\lambda = 0.95$ and set

$$\sigma^2(F_t) = (1 - \delta - \lambda) \times 0.2^2/260 + \delta F_{t-1}^2 + \lambda\sigma^2(F_{t-1}).$$

[8] Instead of assuming a normal distribution for the factor returns and the idiosyncratic components, we assume them to follow t-distributions with 4 degrees of freedom.

[9] So far, the variations change parameters or distributional assumptions within the contagion framework of equation (12). Here we consider another contagion structure by setting up a spillover of volatility instead of a spillover of return shocks: if the idiosyncratic risk of the contagious bank falls below κ, the idiosyncratic risk of the infected banks is m times higher than in the base case. The base level of idiosyncratic risk is chosen such that the total volatility of the infected banks equals the one of the infectious bank. Based on return behavior surrounding the Lehman collapse[20], we assume $m = 3$.

[10] We model a time delay in the spillover. Specifically, we shift 50% of the spillover to the next day:

$$R_{it} = \beta_i F_t + \varepsilon_{it} + 0.5 \sum_{j \neq i} \gamma_j I_{\{\varepsilon_{jt} < \kappa\}} \varepsilon_{jt} + 0.5 \sum_{j \neq i} \gamma_j I_{\{\varepsilon_{j,t-1} < \kappa\}} \varepsilon_{j,t-1}$$

Due to the large number of variations, we cannot present the full set of results for varying contagion intensities. Though the variations may shift the risk measure curves, flatten or elevate them, patterns are mostly very similar to the ones shown in Figure 5. This outcome is exemplified in Table 3, which lists the risk measures for a contagion intensity of 0.75. For this choice, Figure 5 shows that ΔCoVaR assigns a higher risk to the infected banks; exposure ΔCoVaR does as well but with a smaller difference, while MES and beta assign a lower risk to the infected banks. In most variations, the sign of the differences is not changed. For the variations considered here, it never changes for ΔCoVaR, once for MES and twice for beta. Exposure ΔCoVaR shows the largest number of changes, which does not appear surprising given that it showed the most complex behavior in Figure 5.

The findings from Figure 5 therefore seem fairly robust to variations in the parameters and the modeling framework. As in the previous section, they are also robust to distributional assumptions. We find that the risk measures can differ in the way they respond to contagion as it was modeled here. Some risk measures, notably ΔCoVaR, have a tendency to assign a low systemic risk to infectious banks, others tend to do the opposite. One might suspect that the differences between ΔCoVaR measures on the one hand, and MES and beta on the other hand go back to the fact that CoVaR is based on quantiles, and that the findings change when moving to co-expected shortfall (CoES). However, further analyses shows that this is not the case. For example, when we implement the ΔCoES as suggested by Adrian and Brunnermeier (2011) for the base case parameters, the ΔCoES is 1.92% for the infectious bank and 2.58% for the infected bank.

The findings suggest that the use of systemic risk measures could create undesirable incentive effects. Consider again Figure 5, and assume that initially there is no contagion. If a systemic risk charge were based on ΔCoVaR, for example, a bank would have a marginal incentive to become infectious. Its charge would not change, while the charge of its competitors would increase. This would create an advantage for the bank that becomes infectious because its competitors would be required to hold more capital, to pay an insurance premium, or be restricted in some other way. Even if the tighter regulation counteracted the increase in systemic risk that is brought about by an increase in contagion, it would be undesirable to create such an incentive in the first place.

In the base case of Section 3, such effects appear to be most pronounced for ΔCoVaR. However, none of the four measures appears to be immune against creating perverse incentives in the presence of contagion.

[20] For the 29 depositary institutions listed in (Acharya et al. (2012), Appendix A) we examine the idiosyncratic volatility over the 30 days ending on September 12, 2008, as well as over the 30 days starting on September 15, 2008 (Lehman collapse). Using a one-factor model with the S&P 500 as the factor, the median idiosyncratic volatility increases by a factor of 2.98.

Table 3: Systemic risk measures in the presence of contagion: robustness tests.
We simulate returns of N banks. The aggregate system return is the value-weighted average of bank returns. Bank returns are driven by a common factor F, idiosyncratic risk, and bank-to-bank spillovers of idiosyncratic risk:

$$R_i = \beta_i F + \varepsilon_i + \sum_{j \neq i} \gamma_j I_{\{\varepsilon_j < \kappa\}} \varepsilon_j, \; R_S = \sum_{i=1}^{N} w_i R_i.$$

In the base case, parameters (p.a. for drift and volatility) are set to $N = 50$, $\kappa = -0.0204$, $\mathbf{E}(F) = 0.05$, $\sigma_F = 0.2$, $\sigma_i = 0.2$, $\beta_i = 1$, $w_i = 1/50$ for all i. Banks 2 to 50 are not infectious ($\gamma_j = 0$, $\forall j > 1$), while the infectiousness parameter of bank 1 is set to $\gamma_1 = 0.75$. For the variations, the table lists the changes relative to the base case. The measures are estimated through Monte Carlo simulation with 100 million trials. CoVaR_α measures are computed for $\alpha = 1\%$ with observations between the $(\alpha - 0.2\%)$ and $(\alpha + 0.2\%)$ quantiles of the conditioning variable. The MES relates to $\alpha = 5\%$.

	Specification	Bank	ΔCoVaR	Exp. ΔCoVaR	MES	Beta
	Base Case	Infectious	0.0208	0.0315	0.0310	1.0369
		Infected	0.0280	0.0329	0.0294	0.9993
[1]	Equal return mean and	Infectious	0.0209	0.0515	0.0310	1.0392
	volatility across banks	Infected	0.0286	0.0593	0.0284	0.9993
[2]	Threshold $\kappa = -0.0289$	Infectious	0.0208	0.0553	0.0277	1.0112
	$(\mathbf{P}(\varepsilon_1 < \kappa) = 0.01)$	Infected	0.0256	0.0594	0.0270	0.9996
[3]	Threshold $\kappa = -0.0383$	Infectious	0.0207	0.0573	0.0260	1.0020
	$(\mathbf{P}(\varepsilon_1 < \kappa) = 0.001)$	Infected	0.0224	0.0579	0.0259	1.0002
[4]	Factor beta of infectious	Infectious	0.0233	0.0394	0.0363	1.2510
	bank $\beta_1 = 1.25$	Infected	0.0278	0.0333	0.0294	0.9947
[5]	Weight of infectious	Infectious	0.0257	0.0342	0.0337	1.1430
	bank $w_1 = 25\%$	Infected	0.0287	0.0323	0.0290	0.9522
[6]	Five infectious banks	Infectious	0.0373	0.0428	0.0418	1.0240
		Infected	0.0398	0.0436	0.0412	0.9971
[7]	GARCH(1,1) for	Infectious	0.0262	0.0327	0.0318	1.0370
	systematic factor	Infected	0.0303	0.0339	0.0302	0.9991
[8]	t-distribution with 4	Infectious	0.0403	0.0388	0.0355	1.0442
	degrees of freedom	Infected	0.0447	0.0397	0.0331	0.9992
[9]	Volatility spillover instead	Infectious	0.0208	0.0275	0.0259	1.0000
	of return spillover	Infected	0.0237	0.0320	0.0257	1.0003
[10]	50% of spillover is	Infectious	0.0201	0.0335	0.0262	0.9848
	shifted to next day	Infected	0.0227	0.0307	0.0276	1.0001

5 Conclusion

We examine possible pitfalls in the use of return-based measures of systemic risk contributions. Specifically, we check for cases in which a change in an entity's systematic risk, idiosyncratic risk, size or contagiousness increases the risk of the system but lowers the systemic risk contribution of the entity. In such cases, rankings based on estimated systemic risk contributions could produce false interpretations and incentives. In particular, if banks benefit from having a lower estimated systemic contribution in the eyes of their regulators, the use of such measures could motivate banks to take actions that increase the risk of the system rather than reduce it.

While the link between the measured systemic risk contribution and the actual impact on systemic risk is often appropriate, we identify several non-exotic cases in which it is not. In a linear factor model framework with multivariate normal risk factors, we find that the ΔCoVaR measure proposed by Adrian and Brunnermeier (2011) can imply a lower systemic risk contribution if a bank increases its idiosyncratic risk. Beta can lead to unwanted effects if a bank with a large systematic risk increases its size. Exposure ΔCoVaR (Adrian and Brunnermeier (2011)) and marginal expected shortfall (Acharya et al. (2012)) show appropriate sensitivities for all but one situation in which all considered measures behave similarly: if a bank is very dominant, its risk contribution may fall relative to the contribution of another bank if the bank increases its systematic risk.

The last example highlights that a bank can favorably alter its systemic risk contribution relative to those of other banks even though its own risk contribution increases. A study of sensitivities should therefore not be limited to the change in the systemic risk contribution of the bank under analysis but include side effects.

The results are not limited to the return model with multivariate normal risk factors that we start off with. The problematic sensitivities that we identify in this framework also appear when we examine heavy-tailed equity returns generated by a dynamic structural model, or returns drawn from a multivariate t distribution. On top, we find further non-exotic cases of undesired sensitivities that did not appear in the normal model.

Once we introduce contagion into the analysis, we also find differences between the four measures, but none of them is immune against creating false incentives. Stronger spillovers can be associated with a lower systemic risk contribution even though they increase the risk of the system.

Our results should not be interpreted in the sense that measure A is to be preferred to measure B if it shows a lower number of unwanted effects in the analysis. Within our framework, we cannot weigh or compare the importance of specific problems that we identify. In addition, we abstract from estimation error, while systemic risk measures need to be estimated with limited data in practice. Some measures may be less sensitive to estimation error than others.

Despite these caveats, our results are of general relevance and applicability because we show that systemic risk measures can exhibit undesirable properties in standard return frameworks. Knowledge about such limitations is important for practical systemic risk measurement as well as for the development of new measurement approaches.

Appendix A Sensitivities

In this section we explore the sensitivities summarized in Table 1 in a systematic way. We duplicate the table here with references to the analyses provided in this section. As in the main text, a superscript ^{n} marks cases where the opposite sensitivity would appear only under very unusual conditions.

Parameter	Effect type	ΔCoVaR	Exp. ΔCoVaR	MES	Beta
σ_i	direct	$+/-$; A.9	$+$; A.5	$+$; A.5	$+$; A.4
	relative	$+/-$; A.10	$+$; A.11	$+$; A.11	$+$; A.11
β_i	direct	$+$; A.2	$+$; A.6	$+^n$; A.7	$+$; A.6
	relative	$+/-$; A.13	$+/-$; A.12	$+/-$; A.12	$+/-$; A.12
w_i	direct	$+/(-)$; A.3	$+$; A.9	$+$; A.9	$+/-$; A.8
	relative	$+^n$; A.15	$+^n$; A.14	$+^n$; A.14	$+^n$; A.14

We start the analysis with a number of definitions and auxiliary formulas. First we set $w_j^* \equiv (1-w_i)^{-1} w_j$, $j \neq i$, to define bank weights within the sub-system excluding bank i. Corresponding averages of the exposure to systematic risk are $\beta_* \equiv \sum_{k\neq i} w_k^* \beta_k$ and

$$\bar{\beta} \equiv \sum_{k=1}^{N} w_k \beta_k = w_i \beta_i + (1-w_i)\,\beta_* \,.$$

Setting $\varepsilon_* \equiv \sum_{j\neq i} w_j^* \varepsilon_j$ and

$$\bar{\varepsilon} \equiv \sum_{k=1}^{N} w_k \varepsilon_k = w_i \varepsilon_i + (1-w_i)\,\varepsilon_*$$

for aggregate idiosyncratic risks, the index return then reads

$$R_S = \bar{\beta} F + \bar{\varepsilon} = w_i R_i + (1-w_i)\,(\beta_* F + \varepsilon_*)\,,$$

while its variance, using $\sigma_* \equiv \sigma(\varepsilon_*)$, can be written in different ways:

$$\sigma^2(R_S) \;=\; \bar{\beta}^2 \sigma_F^2 + \bar{\sigma}^2 = (w_i\beta_i + (1-w_i)\,\beta_*)^2\,\sigma_F^2 + w_i^2\sigma_i^2 + (1-w_i)^2\,\sigma_*^2 \qquad \text{(A-1)}$$

The covariances $\operatorname{cov}(R_j, R_S)$ play a central role; we denote them by $c_{j,S}$ and obtain the following representation:

$$c_{j,S} \equiv \operatorname{cov}(R_j, R_S) = \beta_j \bar{\beta} \sigma_F^2 + w_j \sigma_j^2 = w_j \sigma^2(R_j) + (1-w_j)\,\beta_j \beta_* \sigma_F^2 \,. \qquad \text{(A-2)}$$

As we assumed all β_j to be positive, all $c_{j,S}$ are positive, too.

A.1 Direct effect of idiosyncratic risk on ΔCoVaR

The new variables turn (3) into

$$\Delta CoVaR_\alpha^{S,i} = \frac{c_{i,S}}{\sigma(R_i)} \Phi^{-1}(1-\alpha) = \left[w_i \sigma(R_i) + (1-w_i)\,\beta_i\beta_* \frac{\sigma_F^2}{\sigma(R_i)} \right] \Phi^{-1}(1-\alpha)\,. \qquad \text{(A-3)}$$

Applying $\partial\sigma(R_i)/\partial\sigma_i = \sigma_i/\sigma(R_i)$, we obtain

$$\begin{aligned} \frac{\partial \Delta CoVaR_\alpha^{S,i}}{\partial \sigma_i} &= \left[w_i \frac{\sigma_i}{\sigma(R_i)} - (1-w_i)\,\beta_i\beta_* \frac{\sigma_F^2}{\sigma^2(R_i)} \frac{\sigma_i}{\sigma(R_i)} \right] \Phi^{-1}(1-\alpha) \\ &\propto \frac{w_i}{1-w_i} - \beta_i\beta_* \frac{\sigma_F^2}{\sigma^2(R_i)} = \frac{w_i}{1-w_i} - \frac{\beta_*}{\beta_i}\left(1 + \frac{\sigma_i^2}{\beta_i^2\sigma_F^2}\right)^{-1}, \end{aligned}$$

which is the formula from which we derived (10) in the special case $\beta_k = 1$ for all k.

A.2 Direct effect of systematic risk on ΔCoVaR

We can directly show that the partial derivative to systematic risk is always positive:

$$\begin{aligned}\frac{\partial}{\partial \beta_i}\left[\Delta CoVaR_{\alpha}^{S,i}\right] &\propto \frac{\partial}{\partial \beta_i}\left[\frac{c_{i,S}}{\sigma\left(R_i\right)}\right] \propto \sigma^2\left(R_i\right)\frac{\partial c_{i,S}}{\partial \beta_i} - c_{i,S}\sigma\left(R_i\right)\frac{\partial \sigma\left(R_i\right)}{\partial \beta_i} \\ &= \sigma^2\left(R_i\right)\frac{\partial c_{i,S}}{\partial \beta_i} - c_{i,S}\frac{1}{2}\frac{\partial \sigma^2\left(R_i\right)}{\partial \beta_i} = \left(\beta_i^2\sigma_F^2+\sigma_i^2\right)\sigma_F^2\left(w_i\beta_i+\bar{\beta}\right) - \left(\beta_i\bar{\beta}\sigma_F^2+w_i\sigma_i^2\right)\beta_i\sigma_F^2 \\ &\propto \left(\beta_i^2\sigma_F^2+\sigma_i^2\right)\left(w_i\beta_i+\bar{\beta}\right) - \left(\beta_i\bar{\beta}\sigma_F^2+w_i\sigma_i^2\right)\beta_i = w_i\beta_i^3\sigma_F^2+\bar{\beta}\sigma_i^2 > 0\,.\end{aligned}$$

A.3 Direct effect of size on ΔCoVaR

We assume that the size of bank i changes while the other banks' size is kept constant. This means that all w_j^* remain constant (and so β_*), which is why (A-3) leads to the simple formula

$$\begin{aligned}\frac{\partial \Delta CoVaR_{\alpha}^{S,i}}{\partial w_i} &= \left[\sigma\left(R_i\right) - \beta_i\beta_*\frac{\sigma_F^2}{\sigma\left(R_i\right)}\right]\Phi^{-1}\left(1-\alpha\right) \\ &\propto -\beta_i\beta_*\sigma_F^2+\sigma^2\left(R_i\right) = \beta_i\,\Delta\beta\,\sigma_F^2+\sigma_i^2\,, \qquad \text{(A-4)}\end{aligned}$$

where $\Delta\beta \equiv \beta_i - \beta_*$. We observe that a bank's weight in the system has an ambiguous effect. Generally, we would expect the ΔCoVaR to increase with size. That is true in many cases, e.g., if $\beta_i \geq \beta_*$. Formula (A-4) shows that the partial derivative is then positive, as it should be.

If, on the contrary, a bank has a rather low exposure to the systematic risk factor and also comparably low idiosyncratic risk, the derivative can have the opposite sign, for instance, if $\sigma_i = 8\%$, $\sigma_F = 20\%$ (both p.a.), $\beta_i = 0.5$, and $\beta_* = 1$.

We show that this is not necessarily a problem because the system return would become less volatile if such a bank gained weight. Assume the partial derivative in (A-4) is negative. It implies $\sigma_i^2 + \Delta\beta\,\beta_i\sigma_F^2 < 0$ and the weaker condition $\Delta\beta < 0$. With $\partial\bar{\beta}/\partial w_i = \bar{\beta}\,\Delta\beta$, we obtain from (A-1):

$$\begin{aligned}\frac{\partial \sigma^2\left(R_S\right)}{\partial w_i} &\propto \bar{\beta}\,\Delta\beta\,\sigma_F^2 + w_i\sigma_i^2 - \left(1-w_i\right)\sigma_*^2 < \left(w_i\beta_i + \left(1-w_i\right)\beta_*\right)\Delta\beta\,\sigma_F^2 + w_i\sigma_i^2 \\ &< w_i\underbrace{\left(\beta_i\Delta\beta\,\sigma_F^2+\sigma_i^2\right)}_{<0} + \left(1-w_i\right)\beta_*\underbrace{\Delta\beta}_{<0}\,\sigma_F^2 < 0\,. \qquad \text{(A-5)}\end{aligned}$$

In short form, we conclude

$$\frac{\partial \Delta CoVaR_{\alpha}^{S,i}}{\partial w_i} < 0 \quad \Rightarrow \quad \frac{\partial \sigma^2\left(R_S\right)}{\partial w_i} < 0.$$

The ΔCoVaR thus sets a correct incentive insofar as a bank is rewarded for growth through a ΔCoVaR-based systemic risk charge only if this lowers the volatility of the system return.

A.4 Direct effect of idiosyncratic risk on beta

Using (A-1) and (A-2), standard calculus shows

$$\begin{aligned}\frac{\partial\, beta_i}{\partial \sigma_i} &\propto \frac{\partial\, beta_i}{\partial \sigma_i^2} = \frac{\partial}{\partial \sigma_i^2}\left[\frac{c_{i,S}}{\sigma^2\left(R_S\right)}\right] \propto \sigma^2\left(R_S\right)\frac{\partial c_{i,S}}{\partial \sigma_i^2} - c_{i,S}\frac{\partial \sigma^2\left(R_S\right)}{\partial \sigma_i^2} \propto \sigma^2\left(R_S\right) - c_{i,S}w_i \\ &= \sum_{k=1}^{N} w_k c_k - c_{i,S}w_i = \sum_{k\neq i} w_k c_k > 0.\end{aligned}$$

A.5 Direct effect of idiosyncratic risk on exposure ΔCoVaR and MES

$\sigma(R_S)$ is obviously a growing function of σ_i^2, and so is $beta_i$. The exposure ΔCoVaR equals $\Phi^{-1}(1-\alpha) \times beta_i \times \sigma(R_S)$, which is a product of two monotonic functions of σ_i^2 and a constant. Hence, the exposure ΔCoVaR is increasing in σ_i.

As the MES differs from the exposure ΔCoVaR only by a constant factor and an offset that does not depend on σ_i, it shows the same monotonicity w.r.t. idiosyncratic risk.

A.6 Direct effect of systematic risk on beta and exposure ΔCoVaR

We first consider beta:

$$\begin{aligned} \frac{\partial\, beta_i}{\partial \beta_i} &= \frac{\partial}{\partial \beta_i}\left[\frac{c_{i,S}}{\sigma^2(R_S)}\right] \propto \sigma^2(R_S)\frac{\partial c_{i,S}}{\partial \beta_i} - c_{i,S}\frac{\partial \sigma^2(R_S)}{\partial \beta_i} \\ &= \left(\bar{\beta}^2\sigma_F^2 + \bar{\sigma}^2\right)\left[\bar{\beta} + w_i\beta_i\right]\sigma_F^2 - \left(\beta_i\bar{\beta}\sigma_F^2 + w_i\sigma_i^2\right)2\bar{\beta}w_i\sigma_F^2 \,. \end{aligned}$$

We divide by σ_F^2 and estimate the outcome from below by removing the second addend of $\bar{\sigma}^2 = w_i^2\sigma_i^2 + (1-w_i)^2\sigma_*^2$ so that we obtain

$$\begin{aligned} \frac{\partial\, beta_i}{\partial \beta_i} &\propto \ldots \geq \left(\bar{\beta}^2\sigma_F^2 + w_i^2\sigma_i^2\right)\left[\bar{\beta} + w_i\beta_i\right] - 2\bar{\beta}w_i\beta_i\bar{\beta}\sigma_F^2 - 2\bar{\beta}w_iw_i\sigma_i^2 \\ &= \bar{\beta}^3\sigma_F^2 + w_i\beta_i\bar{\beta}^2\sigma_F^2 + w_i^2\bar{\beta}\sigma_i^2 + w_i^3\beta_i\sigma_i^2 - 2w_i\beta_i\bar{\beta}^2\sigma_F^2 - 2w_i^2\bar{\beta}\sigma_i^2 \\ &= \bar{\beta}^3\sigma_F^2 + w_i^3\beta_i\sigma_i^2 - w_i\beta_i\bar{\beta}^2\sigma_F^2 - w_i^2\bar{\beta}\sigma_i^2 \\ &= \bar{\beta}^2\sigma_F^2\bar{\beta} + w_i^3\beta_i\sigma_i^2 - w_i\beta_i\bar{\beta}^2\sigma_F^2 - w_i\beta_iw_i^2\sigma_i^2 - (1-w_i)\beta_*w_i^2\sigma_i^2 \\ &= (1-w_i)\beta_*\left(\bar{\beta}^2\sigma_F^2 - w_i^2\sigma_i^2\right) \propto \bar{\beta}^2\sigma_F^2 - w_i^2\sigma_i^2 \,. \end{aligned}$$

Hence, we can state

$$w_i < \frac{\bar{\beta}\sigma_F}{\sigma_i} \quad \Rightarrow \quad \frac{\partial\, beta_i}{\partial \beta_i} > 0 \,. \tag{A-6}$$

The condition is fulfilled unless a bank is either extremely dominant in the system or has huge idiosyncratic risk. For example, assuming that a bank's idiosyncratic risk does not exceed the double of all banks' average systematic risk, as given by $\bar{\beta}\sigma_F$, this bank would need to make up more than half of the banking system to make the beta fall in β_i.

As regards the exposure ΔCoVaR, we apply the same argument as in Appendix A.5: the exposure ΔCoVaR, $\Phi^{-1}(\alpha) \times beta_i \times \sigma(R_S)$, is an increasing function of β_i because its components $\sigma(R_S)$ and $beta_i$ are increasing.

A.7 Direct effect of systematic risk on MES

Recall $MES_i = -\beta_i\mu + beta_i\,\sigma(R_S)\,C_\alpha$, where $C_\alpha \equiv \alpha^{-1}\phi\left(\Phi^{-1}(\alpha)\right)$ is a positive constant. The MES is special because of its drift related term, which decreases in β_i. The partial derivative is

$$\frac{\partial MES_i}{\partial \beta_i} = -\mu + C_\alpha \frac{\partial}{\partial \beta_i}\left[beta_i\,\sigma(R_S)\right], \tag{A-7}$$

of which we know from (A-6) that the partial derivative on the right-hand side is positive. If we can neglect μ, we know that the MES is an increasing function of β_i. Whether it actually can be neglected depends on the risk horizon. While μ is proportional to the risk horizon, $\sigma(R_S)$ is so to its square root. As the risk horizonis 1 day throughout this paper , and since C_α is already larger than 2 for $\alpha = 5\%$, the first term is by magnitudes smaller than the second. The MES will therefore be an increasing function of systematic risk under all plausible conditions. Even on an annual basis, $-\mu$ would usually be dominated by the positive part in (A-7).[21]

A.8 Direct effect of size on beta

As the partial derivative of $beta_i$ is complicated, we make the simplifying assumption that only bank i has a non-negligible weight in the system, whereas all other banks are infinitesimally small. In the limit, idiosyncratic risks of these banks are diversified away ($\sigma_* = 0$), so that we obtain:

$$\begin{aligned}\frac{\partial\, beta_i}{\partial w_i} &= \frac{\partial}{\partial w_i}\left[\frac{c_{i,S}}{\sigma^2(R_S)}\right] \propto \sigma^2(R_S)\frac{\partial c_{i,S}}{\partial w_i} - c_{i,S}\frac{\partial \sigma^2(R_S)}{\partial w_i} = \sigma^2(R_S)\frac{\partial c_{i,S}}{\partial w_i} - c_{i,S}\frac{\partial \sigma^2(R_S)}{\partial w_i} \\ &= \left(\bar\beta^2\sigma_F^2 + w_i^2\sigma_i^2\right)\left(\beta_i\,\Delta\beta\,\sigma_F^2 + \sigma_i^2\right) - \left(\beta_i\bar\beta\sigma_F^2 + w_i\sigma_i^2\right)2\left(\bar\beta\,\Delta\beta\sigma_F^2 + w_i\sigma_i^2\right),\end{aligned}$$

where $\Delta\beta \equiv \beta_i - \beta_*$. As the conditions under which this expression is positive are still difficult to identify, we focus on the case where w_i is also small enough to set it zero. We obtain

$$\left.\frac{\partial\, beta_i}{\partial w_i}\right|_{w_i=0} \propto \sigma_i^2 - \beta_i\,\Delta\beta\,\sigma_F^2\,. \tag{A-8}$$

This derivative should preferably be positive. It is not, however, if the bank's exposure to systematic risk is above the average and the idiosyncratic risk is comparably small. Note that the growth of such a bank would increase the variance of the system return, which can be seen in (A-5), where we have found

$$\frac{\partial\sigma^2(R_S)}{\partial w_i} \propto \bar\beta\,\Delta\beta\,\sigma_F^2 + w_i\sigma_i^2 - (1-w_i)\,\sigma_*^2\,. \tag{A-9}$$

Under the limiting assumption $\sigma_* = 0$, the variance of R_S grows in w_i if $\beta_i > \beta_*$.

A.9 Direct effect of size on exposure ΔCoVaR and MES

We first consider the main part of the exposure ΔCoVaR, $c_{i,S}/\sigma(R_S)$. Using the property $\partial\sigma(R_S)/\partial w_i = (2\sigma(R_S))^{-1}\,\partial\sigma^2(R_S)/\partial w_i$, we find:

$$\begin{aligned}\frac{\partial}{\partial w_i}\left[\frac{c_{i,S}}{\sigma(R_S)}\right] &\propto \sigma^2(R_S)\frac{\partial c_{i,S}}{\partial w_i} - \frac{1}{2}c_{i,S}\frac{\partial\sigma^2(R_S)}{\partial w_i} \\ &= \left(\beta_i\,\Delta\beta\,\sigma_F^2 + \sigma_i^2\right)\left(\bar\beta^2\sigma_F^2 + w_i^2\sigma_i^2 + (1-w_i)^2\sigma_*^2\right) \\ &\quad - \left(\beta_i\bar\beta\sigma_F^2 + w_i\sigma_i^2\right)\left(\bar\beta\,\Delta\beta\sigma_F^2 + w_i\sigma_i^2 - (1-w_i)\,\sigma_*^2\right) \\ &= \sigma_F^2\sigma_i^2\left\{\beta_i\Delta\beta\, w_i^2 - \beta_i\bar\beta w_i - \bar\beta\,\Delta\beta\, w_i + \bar\beta^2\right\} + \sigma_i^2\sigma_*^2(1-w_i) \\ &\quad + \sigma_F^2\sigma_*^2\beta_i(1-w_i)\left[\Delta\beta(1-w_i) + \bar\beta\right] \\ &= \ldots = \sigma_F^2\sigma_i^2(1-w_i)\,\beta_*^2 + \sigma_i^2\sigma_*^2(1-w_i) + \sigma_F^2\sigma_*^2\beta_i(1-w_i)(1-w_i)\,\beta_i > 0\end{aligned}$$

This means that the exposure ΔCoVaR is always an increasing function of w_i. The same holds for the MES since the additive term $-\mu\beta_i$ in (7) is independent of w_i.

[21] Assuming $\alpha \leq 5\%$, the drift over one year would have to be roughly twice as large as the annual market volatility to make the MES fall in β_i. This would be a very unusual market. Details of the corresponding estimate are available on request.

A.10 Relative effect of idiosyncratic risk on ΔCoVaR

The parameter σ_i has neither an impact on another bank's return volatility $\sigma(R_j)$ nor on its covariance $c_{j,S}$ with the system return. According to $\Delta CoVaR_\alpha^{S,j} = c_{j,S}/\sigma(R_j)\,\Phi^{-1}(1-\alpha)$, the ΔCoVaR of bank j is then also unaffected. Hence, the relative and the direct effect of σ_i fall together, apart from a constant factor.

A.11 Relative effect of idiosyncratic risk on exposure ΔCoVaR, MES, and beta

We consider the ratio of two banks' systemic risk measures, such as $beta_i/beta_j$. For exposure ΔCoVaR and beta the ratios are equal to $c_{i,S}/c_{j,S}$. As already stated, $c_{j,S}$ is invariant to σ_i so that only the effect on the covariance $c_{i,S}$ remains to be analyzed. It is obviously positive because $\partial c_{i,S}/\partial \sigma_i = 2w_i\sigma_i > 0$. The MES has drift related addends above and below the fraction line. We neglect them in this section as they are small, based on the arguments provided in Appendix A.7. We therefore consider the MES to be covered by the analysis of $c_{i,S}/c_{j,S}$.

A.12 Relative effect of systematic risk on exposure ΔCoVaR, MES, and beta

For the partial derivative of the ratio of covariances to β_i, we obtain:

$$\begin{aligned}\frac{\partial}{\partial\beta_i}\left[\frac{c_{i,S}}{c_{j,S}}\right] &\propto c_{j,S}\frac{\partial c_{i,S}}{\partial\beta_i} - c_{i,S}\frac{\partial c_{j,S}}{\partial\beta_i} = \left(\beta_j\bar{\beta}\sigma_F^2 + w_j\sigma_j^2\right)\sigma_F^2\left(w_i\beta_i+\bar{\beta}\right) - \left(\beta_i\bar{\beta}\sigma_F^2 + w_i\sigma_i^2\right)\sigma_F^2\beta_j w_i \\ &\propto 1 + w_j\frac{\sigma_j^2}{\beta_j\bar{\beta}^2\sigma_F^2}\left(w_i\beta_i+\bar{\beta}\right) - w_i^2\frac{\sigma_i^2}{\bar{\beta}^2\sigma_F^2}\,. \end{aligned} \tag{A-10}$$

In absence of a dominating bank, the only negative part of the expression is considerably smaller than 1 under most conditions because of the factor w_i^2. The relative effect of β_i is then positive. However, it may become negative if bank i is really large. Assume for simplicity that bank j is very small so that the middle term vanishes. Then, the ratio $c_{i,S}/c_{j,S}$ will negatively depend on β_i if $w_i\sigma_i > \bar{\beta}\sigma_F$.

A.13 Relative effect of systematic risk on ΔCoVaR

We start with a calculation of the partial derivative

$$\begin{aligned}\frac{\partial}{\partial\beta_i}\left[\frac{\Delta CoVaR_\alpha^{S,i}}{\Delta CoVaR_\alpha^{S,j}}\right] &= \sigma(R_j)\frac{\partial}{\partial\beta_i}\left[\frac{c_{i,S}}{c_{j,S}\sigma(R_i)}\right] \propto \sigma(R_i)\frac{\partial}{\partial\beta_i}\left[\frac{c_{i,S}}{c_{j,S}}\right] - \frac{c_{i,S}}{c_{j,S}}\frac{\partial\sigma(R_i)}{\partial\beta_i} \\ &= \sigma(R_i)\frac{\sigma_F^2}{c_{j,S}}\left\{\beta_j\bar{\beta}^2\sigma_F^2 + w_j\sigma_j^2\bar{\beta} + \left(w_j\sigma_j^2\beta_i - w_i\beta_j\sigma_i^2\right)w_i\right\} - \frac{c_{i,S}}{c_{j,S}}\frac{\sigma_F^2}{\sigma(R_i)}\beta_i \\ &\propto \sigma^2(R_i)\left\{\beta_j\bar{\beta}^2\sigma_F^2 + w_j\sigma_j^2\bar{\beta} + \left(w_j\sigma_j^2\beta_i - w_i\beta_j\sigma_i^2\right)w_i\right\} - c_{i,S}c_{j,S}\beta_i\,. \end{aligned}$$

It is difficult to determine under which conditions this expression becomes negative. We therefore use an approximation where we assume that bank i may dominate the system whereas the weight of the benchmark bank j can be neglected.[22] Eliminating all terms containing w_j gives

$$\begin{aligned}
\frac{\partial}{\partial \beta_i}\left[\frac{\Delta CoVaR_\alpha^{S,i}}{\Delta CoVaR_\alpha^{S,j}}\right] &\propto \ldots \approx \sigma^2\left(R_i\right)\left\{\beta_j\bar{\beta}^2\sigma_F^2 - w_i^2\beta_j\sigma_i^2\right\} - c_{i,S}c_{j,S}\beta_i \\
&= \left(\beta_i^2\sigma_F^2+\sigma_i^2\right)\left\{\beta_j\bar{\beta}^2\sigma_F^2 - w_i^2\beta_j\sigma_i^2\right\} - \beta_i\left(\beta_i\bar{\beta}\sigma_F^2+w_i\sigma_i^2\right)\left(\beta_j\bar{\beta}\sigma_F^2+w_j\sigma_j^2\right) \\
&\approx \left(\beta_i^2\sigma_F^2+\sigma_i^2\right)\left\{\beta_j\bar{\beta}^2\sigma_F^2 - w_i^2\beta_j\sigma_i^2\right\} - \left(\beta_i\bar{\beta}\sigma_F^2+w_i\sigma_i^2\right)\beta_i\beta_j\bar{\beta}\sigma_F^2\,.
\end{aligned}$$

Further consolidation plus introduction of $\kappa \equiv w_i/\left(1-w_i\right)$ and $\rho \equiv \beta_i/\beta_*$ lead to

$$\begin{aligned}
\frac{\partial}{\partial \beta_i}\left[\frac{\Delta CoVaR_\alpha^{S,i}}{\Delta CoVaR_\alpha^{S,j}}\right] &\approx \ldots \propto \sigma_F^2\left(\beta_i^2\bar{\beta}^2\sigma_F^2 - \beta_i^2 w_i^2\sigma_i^2 + \bar{\beta}^2\sigma_i^2 - w_i^2\sigma_i^4 - \beta_i^2\bar{\beta}^2\beta_j\sigma_F^2 - \beta_i w_i\sigma_i^2\bar{\beta}\right) \\
&\propto \bar{\beta}^2\sigma_F^2 - w_i^2\left(\sigma_i^2+\beta_i^2\sigma_F^2\right) - \beta_i w_i\bar{\beta}\sigma_F^2 = \sigma_F^2\bar{\beta}\left[\bar{\beta}-\beta_i w_i\right] - w_i^2\left(\sigma_i^2+\beta_i^2\sigma_F^2\right) \\
&= \left(1-w_i\right)\sigma_F^2\bar{\beta}\beta_* - w_i^2\left(\sigma_i^2+\beta_i^2\sigma_F^2\right) = \left(1-w_i\right)\left[w_i\beta_i+\left(1-w_i\right)\beta_*\right]\beta_* - w_i^2\left(\frac{\sigma_i^2}{\sigma_F^2}+\beta_i^2\right) \\
&\propto \left[\kappa\beta_i+\beta_*\right]\beta_* - \kappa^2\left(\frac{\sigma_i^2}{\sigma_F^2}+\beta_i^2\right) = \beta_*^2 - \kappa\left(\kappa\frac{\sigma_i^2}{\sigma_F^2}+\beta_i\left(\kappa\beta_i-\beta_*\right)\right) \\
&\propto 1-\kappa\left(\kappa\frac{\sigma_i^2}{\sigma_F^2\beta_*^2}+\rho\left(\kappa\rho-1\right)\right).
\end{aligned}$$

This expression can be positive or negative; it is discussed in the main text above Equation (11).

A.14 Relative effect of size on MES, exposure ΔCoVaR, and beta

The partial derivative of the ratio of covariances can be simplified to:

$$\begin{aligned}
\frac{\partial}{\partial w_i}\left[\frac{c_{i,S}}{c_{j,S}}\right] &\propto c_{j,S}\frac{\partial c_{i,S}}{\partial w_i} - c_{i,S}\frac{\partial c_{j,S}}{\partial w_i} \\
&= \left(\beta_j\bar{\beta}\sigma_F^2+w_j\sigma_j^2\right)\left(\sigma_i^2+\beta_i\,\Delta\beta\,\sigma_F^2\right) - \left(\beta_i\bar{\beta}\sigma_F^2+w_i\sigma_i^2\right)\beta_j\Delta\beta\sigma_F^2 \\
&= \left(\beta_j\sigma_F^2 w_i\beta_i+\left(1-w_i\right)\beta_j\beta_*\sigma_F^2+w_j\sigma_j^2\right)\sigma_i^2 + w_j\sigma_j^2\beta_i\,\Delta\beta\,\sigma_F^2 - w_i\sigma_i^2\beta_j\,\Delta\beta\,\sigma_F^2 \\
&= \left(\beta_j\beta_*\sigma_F^2+w_j\sigma_j^2\right)\sigma_i^2 + w_j\sigma_j^2\beta_i\,\Delta\beta\,\sigma_F^2 = \beta_j\beta_*\sigma_F^2\sigma_i^2 + w_j\sigma_j^2\left(\sigma_i^2+\beta_i\,\Delta\beta\,\sigma_F^2\right) \qquad \text{(A-11)}
\end{aligned}$$

The derivative can actually become negative but only if the system volatility is a falling function of w_i. In fact, inspecting the last line, the derivative can only be negative if $\sigma_i^2+\beta_i\Delta\beta\sigma_F^2<0$; this condition is sufficient for $\partial\sigma^2\left(R_S\right)/\partial w_i$, as shown in (A-5). The case is similar to the direct effect of size on the ΔCoVaR (Appendix A.3), which we classified to be ambiguous but, in a sense, appropriate. However, the conditions under which the partial derivative in (A-11) can become negative are considerably more exotic, as may be illustrated by the following estimate (details omitted):

$$\frac{\partial}{\partial w_i}\left[\frac{c_{i,S}}{c_{j,S}}\right] \propto \ldots \geq 1-\frac{1}{4}w_j\frac{\sigma_j^2\beta_*}{\sigma_i^2\beta_j}.$$

A.15 Relative effect of size on ΔCoVaR

The effect can be traced back to (A-11) since $\sigma\left(R_i\right)$ and $\sigma\left(R_j\right)$ are invariant to w_i:

$$\frac{\partial}{\partial w_i}\left[\frac{\Delta CoVaR_\alpha^{S,i}}{\Delta CoVaR_\alpha^{S,j}}\right] = \frac{\sigma\left(R_j\right)}{\sigma\left(R_i\right)}\frac{\partial}{\partial w_i}\left[\frac{c_{i,S}}{c_{j,S}}\right] \propto \frac{\partial}{\partial w_i}\left[\frac{c_{i,S}}{c_{j,S}}\right].$$

[22] The relative effect in the reversed case (where bank i is small) is likely to be very similar to the direct effect since the small bank has only weak impact on the index, and so is its effect on the large bank's ΔCoVaR. Hence, the small bank's relative effect is basically the effect on its own ΔCoVaR, divided by a constant.

The relative effect is then the same as for the other measures.

Appendix B Distribution of the noise term ε_1, conditional on the return

Given the return $R_1 = \beta_1 F + \varepsilon_1$ of the infectious bank in Section 4, we calculate the distribution of ε_1 under the condition that R_1 equals its quantile value at level α. This is done by an orthogonal linear representation, in the same way as in Section 3. Starting from the assumptions $F \sim N\left(\mu, \sigma_F^2\right)$, $\varepsilon_1 \sim N\left(0, \sigma_1^2\right)$ and independence of F and ε_1, we calibrate an equation $\varepsilon_1 = a + bR_1 + \eta$ such that R_1 and η are independent and $\mathbf{E}\varepsilon_1 = \mathbf{E}\eta = 0$. This requires $a = -b\beta_1\mu$ and

$$\varepsilon_1 = b\left(R_1 - \beta_1\mu\right) + \eta \tag{B-1}$$

with $b = \operatorname{cov}\left(R_1, \varepsilon_1\right)/\sigma^2\left(R_1\right) = \sigma_1^2/\sigma^2\left(R_1\right)$. The noise orthogonal to R_1 has variance

$$\sigma_\eta^2 = \sigma_1^2 - b^2\sigma^2\left(R_1\right) = \sigma_1^2\beta_1^2\sigma_F^2/\sigma^2\left(R_1\right) ,$$

which completes the conditional moments. Putting $Q_\alpha\left(R_1\right) = \beta_1\mu + \sigma\left(R_1\right)\Phi^{-1}\left(\alpha\right)$ into (B-1), we obtain

$$\varepsilon_1 \left|\left\{R_1 = Q_\alpha\left(R_1\right)\right\}\right. \sim N\left(b\left(Q_\alpha\left(R_1\right) - \beta_1\mu\right), \sigma_\eta^2\right) = N\left(\frac{\sigma_1^2}{\sigma\left(R_1\right)}\Phi^{-1}\left(\alpha\right), \frac{\sigma_1^2\beta_1^2\sigma_F^2}{\sigma^2\left(R_1\right)}\right). \tag{B-2}$$

We are also interested in the conditional probability of the contagion event, which is $\{\varepsilon_1 < \kappa\}$. Let us assume that the unconditional probability of contagion is χ. The R_1-conditional probability is calculated by standardization of ε_1 based on the moments given in (B-2):

$$\begin{aligned}\mathbf{P}\left(\varepsilon_1 < \kappa | R_1 = Q_\alpha\left(R_1\right)\right) &= \mathbf{P}\left(\varepsilon_1 < \sigma_1\Phi^{-1}\left(\chi\right) | R_1 = Q_\alpha\left(R_1\right)\right) \\ &= \Phi\left(\frac{\sigma_1\Phi^{-1}\left(\chi\right) - \mathbf{E}\left(\varepsilon_1 | R_1 = Q_\alpha\left(R_1\right)\right)}{\sigma\left(\varepsilon_1 | R_1 = Q_\alpha\left(R_1\right)\right)}\right) = \Phi\left(\frac{\sigma\left(R_1\right)\Phi^{-1}\left(\chi\right) - \sigma_1\Phi^{-1}\left(\alpha\right)}{\beta_1\sigma_F}\right).\end{aligned} \tag{B-3}$$

Appendix C The structural model for asset and equity returns

In this section we present the structural model used in the robustness test for the linear case. We extend (and simplify) the model of Collin-Dufresne and Goldstein (2001), which has been selected since it is one of the few that generate stationary returns both for assets and equity.[23]

We model asset returns as in the linear normal model of Section 3, with the modification that returns over finite time intervals are now lognormal. The SDE system for the latent systematic factor F_t and asset values $V_{i,t}$ reads

$$\frac{\mathrm{d}F_t}{F_t} = \mu\,\mathrm{d}t + \sigma_F\,\mathrm{d}B_t, \quad \frac{\mathrm{d}V_{i,t}}{V_{i,t}} = \beta_i\frac{\mathrm{d}F_t}{F_t} + \sigma_i\,\mathrm{d}B_{i,t}$$

with independent Brownian motions B_t and $B_{i,t}$. The asset returns are stationary by construction. It is convenient to replace the independent Brownian motions by the N-dimensional Gaussian process

$$Z_t \equiv \left(\beta_i\sigma_F B_{F,t} + \sigma_i B_{i,t}\right)_{i=1}^N ,$$

[23] The models of Leland (1994) and Leland and Toft (1996) might appear as natural alternatives since they include stationary debt pricing. However, neither equity returns nor those of the market value of assets are stationary.

which has zero drift and the covariance function

$$\Omega_{ij}(t,s) \equiv \text{cov}(Z_{i,t}, Z_{j,s}) = \left(\beta_i \beta_j \sigma_F^2 + \sigma_i^2 I_{\{i=j\}}\right) \min(t,s) \,. \tag{C-1}$$

Most calculations are done in logarithmic terms, which we denote by small characters. The log assets process $v_{i,t} \equiv \log(V_{i,t})$ of bank i is an arithmetic Brownian motion following

$$\mathrm{d}v_{i,t} = \eta_i \,\mathrm{d}t + \mathrm{d}Z_{i,t} \quad \text{with} \quad \eta_i \equiv \beta_i \mu - 0.5\left(\beta_i^2 \sigma_F^2 + \sigma_i^2\right) \,. \tag{C-2}$$

Each bank steers its debt by corporate action in order to achieve a certain target leverage.[24] The model approximates this behavior by a controlled dynamic default threshold $K_{i,t}$, which we interpret as the balance sheet value of debt. It is time-differentiable and assumed to follow, in its logarithmic form, the ODE

$$\mathrm{d}k_{i,t} = \left[\lambda_i \left(v_{i,t} - k_{i,t} + \bar{l}_i\right) + \beta_i \mu\right] \mathrm{d}t \,,$$

where target leverage $\bar{l}_i$ and adjustment speed λ_i are strategic parameters.[25] Logarithmic "leverage" is defined as the distance $l_{i,t} = k_{i,t} - v_{i,t}$ between the log default threshold and log assets. As long as the bank is alive, $l_{i,t}$ is an Ornstein-Uhlenbeck (OU) process

$$\mathrm{d}l_{i,t} = \lambda_i \left(\bar{l}_i - l_{i,t}\right) \mathrm{d}t - \mathrm{d}Z_{i,t} \,.$$

Normally, default would occur at the first time when $K_{i,t} = V_{i,t}$ holds or, equivalently, $l_{i,t} = 0$. As we are interested in observable time series for banks, and for technical reasons, we assume that the supervisor would take a bank into conservatorship and remove it from stock markets when its equity, relative to assets, falls short of a small but positive amount. Formally, we define the "default" time as $\tau_i \equiv \inf\{t : l_{i,t} = l_{\max}\}$ and set $l_{\max} = \log(97\%)$ in the simulations.

Equity is defined as the difference between the market values of assets and debt. For simplicity, we assume the accounting and market value of debt to be identical, so that equity is just $E_{i,t} = V_{i,t} - K_{i,t}$.[26] Normally, a shock to $V_{i,t}$ would partly carry over to the market value of $K_{i,t}$, and especially so over short-term horizons where the adaptation of the smooth process $K_{i,t}$ is of second order, compared to the diffusion shocks to $V_{i,t}$. In our simplified model, however, the short-term variation of assets *completely* carries over to the value of equity, which makes it more volatile especially in moments of high leverage, compared to a model with precise debt pricing. As we mainly test whether our results are robust to the presence of heavier tails in return distributions, we find it acceptable that these tails are a bit heavier than those arising from a fully-fledged debt pricing model.

Using

$$\mathrm{d}E_{i,t} = \mathrm{d}V_{i,t} - \mathrm{d}K_{i,t} = \left[\beta_i \mu E_{i,t} + \lambda_i K_{i,t}\left(l_{i,t} - \bar{l}_i\right)\right] \mathrm{d}t + V_{i,t} \mathrm{d}Z_{i,t} \,,$$

Itô's lemma gives the following SDE for log equity:

$$\mathrm{d}\log E_{i,t} = \left[\beta_i \mu + \lambda_i \frac{l_{i,t} - \bar{l}_i}{e^{-l_{i,t}} - 1} - \frac{1}{2} \frac{\beta_i^2 \sigma_F^2 + \sigma_i^2}{\left(1 - e^{l_{i,t}}\right)^2}\right] \mathrm{d}t + \frac{1}{1 - e^{l_{i,t}}} \mathrm{d}Z_{i,t} \tag{C-3}$$

[24] Such corporate action can have various forms but is most conveniently thought of as purely liabilities-related transactions, such as debt/equity swaps or debt-financed stock repurchases.

[25] By adding $\beta_i \mu$ to the drift of $\kappa_{i,t}$, we differ from Collin-Dufresne and Goldstein (2001) in that our parameter $\bar{l}_i$ actually equals the expectation of $l_{i,t}$ under the stationary measure; in the original work there is a gap between them. The difference in the parameters is only a matter of notation.

[26] This assumption is not found in the work of Collin-Dufresne and Goldstein (2001). Focusing on bond pricing, they do not need to model the value of debt and equity explicitly. The only link between their structural model and bond pricing is the distribution of the default time, which is already defined by the Ornstein-Uhlenbeck process.

The formula shows two things. First, as long as the log leverage $l_{i,t}$ is stationary, equity returns are stationary, too. Second, the dynamic diffusion generates heteroskedasticity in the equity returns.[27]

We now write target leverages and adjustment speeds in vector form $\bar{l}$ and λ and specify the stationary distribution of $l_t = (l_{i,t})_{i=1}^N$. If we could ignore that processes are stopped at τ_i, the stationary distribution would be $N\left(\bar{l}, \Sigma\right)$, where $\Sigma_{ij} = \left(\beta_i\beta_j\sigma_F^2 + \sigma_i^2 I_{\{i=j\}}\right) / \left(\lambda_i + \lambda_j\right)$. Of course, stopping cannot be ignored since, otherwise, some of the processes would have to start in the default state. We therefore select a distribution of l_t that is stationary *conditional on survival*, meaning that it fulfills

$$\mathbf{P}\left(l_s \in B | \tau_i > s,\ i = 1, ..., N\right) = \mathbf{P}\left(l_t \in B | \tau_i > t,\ i = 1, ..., N\right)$$

for arbitrary times t, s and measurable sets $B \subset \mathbb{R}^N$. This distribution is not analytically available; we approximate it by simulation as described below.

For simulation purposes, we replace the SDE (C-3) by an equation where drift and volatility are kept constant in a small time interval, in our case one day. The simulation of one-day asset and equity returns consists of the following steps. As explained below, multiple independent simulations must be performed in parallel.

1. Seed sample: draw M independent instances from a truncation of the multivariate $N\left(\bar{l}, \Sigma\right)$ distribution to the set $\left(-\infty, l_{\max}\right)^N$, where $l_{\max}$ is the uniform stopping threshold for $l_{i,t}$.

2. Draw M independent instances of the one-day diffusion term from a normal distribution: $\varepsilon \sim N\left(0, \Omega\left(T, T\right)\right)$, where $T \equiv 1/260$ is one trading day and Ω is defined in (C-1). Log leverage of the next day is obtained from[28] $l_1 = l_0 + diag(\lambda)\left(\bar{l} - l_0\right) - \varepsilon$ which, however, can also end up with some values larger than $l_{\max}$. As we censor stopping events, in such a case the l_0 is replaced by a randomly selected instance of l_0 from the sample, and a new ε is drawn. If necessary, the replacement is repeated until l_1 is smaller than $l_{\max}$ in all components.[29]

3. Having obtained l_1 from step 2, set $l_0 \equiv l_1$ and go back to step 2. Repeat this loop until the distribution converges to survival-conditional stationarity.[30] After convergence, go to the next step.

4. Draw $\varepsilon \sim N\left(0, \Omega\left(T, T\right)\right)$. Apply (C-2) to calculate daily asset returns as

$$R_{i,V} \equiv \exp\left\{\eta_i T + \varepsilon_i\right\} - 1\,.$$

5. Randomly pick one of the M instances of l_0. Calculate one-day equity returns, according to (C-3), but keeping coefficients constant for one day, as

$$R_{i,eq} \quad \equiv \quad \left[\beta_i\mu + \lambda_i \frac{l_{i,0} - \bar{l}_i}{e^{-l_{i,0}} - 1} - \frac{1}{2}\frac{\beta_i^2\sigma_F^2 + \sigma_i^2}{\left(1 - e^{l_{i,0}}\right)^2}\right] T + \frac{1}{1 - e^{l_{i,0}}}\varepsilon_i\,.$$

[27] The solution of the SDE could explode if we allowed $l_{i,t}$ to reach zero. To prevent technical problems, we stop the process at τ_i, which bounds the diffusion differential from above. In our simulations, the instantaneous equity volatility can, at max, be about 3.3 times larger than the average.

[28] This AR(1) process is an approximation of the Ornstein-Uhlenbeck process. We could also set mean reversion and variance of the AR process such that it has exactly the same distribution as the Ornstein-Uhlenbeck process observed at discrete times; however, these parameters are almost exactly the same as λ and ε.

[29] If we knew the stationary distribution in advance, resetting l_0 would not be necessary. It is so to achieve convergence to the survival-conditional stationary distribution.

[30] We test for survival-conditional stationarity by the convergence of the sample characteristics mean, variance, skewness and kurtosis.

6. Add R_V and R_{eq} to the sample and go back to step 4.[31]

As in the base case, the index return is defined as a weighted average. For the simulations we set the following base case parameters. $N = 50$ homogeneous banks of equal size; $\sigma_F = 0.05$, $\mu = 0.03$, $\sigma_i = 0.04$ (all annualized); $\beta_i = 1$, $\bar{l}_i = -0.1$; $\lambda_i = 2.38$. Note that the values relate to bank asset returns, which are typically much less volatile than those of corporates. The risk parameters are roughly consistent with KMV asset volatilities of banks and the values found by Memmel and Raupach (2010).[32] We generate $M = 50,000$ samples for l_0, based on a sequence of 100 days to achieve survival-conditional stationarity. Estimates of systemic risk measures are based on 10 million return vectors.

References

Acharya, V., R. Engle, and M. Richardson (2012). Capital Shortfall: A New Approach to Ranking and Regulating Systemic Risks. *American Economic Review 102*(3), 59–64.

Acharya, V. V., L. H. Pedersen, T. Philippon, and M. P. Richardson (2012). Measuring systemic risk. CEPR Discussion Paper 8824.

Adams, Z., R. Füss, and R. Gropp (2014). Spillover effects among financial institutions: A state-dependent sensitivity value-at-risk approach. *Journal of Financial and Quantitative Analysis 49*, 575–598.

Adrian, T. and M. K. Brunnermeier (2011). CoVaR. NBER Working Paper w17454.

Allen, F. and E. Carletti (2013). What is systemic risk? *Journal of Money, Credit and Banking 45*, 121–127.

Artzner, P., F. Delbaen, J. Eber, and D. Heath (1999). Coherent measures of risk. *Mathematical Finance 9*(3), 203–228.

Baele, L., O. De Jonghe, and R. Vander Vennet (2007). Does the Stock Market Value Bank Diversification? *Journal of Banking & Finance 31*(7), 1999–2023.

Barth, A. and I. Schnabel (2013). Why banks are not too big to fail – evidence from the CDS market. *Economic Policy 28*(74), 335–369.

Basak, S. and A. Shapiro (2001). Value-at-risk-based risk management: Optimal policies and asset prices. *Review of Financial Studies 14*(2), 371–405.

Battaglia, F. and A. Gallo (2013). Securitization and systemic risk: An empirical investigation on Italian banks over the financial crisis. *International Review of Financial Analysis 30*, 274–286.

Benoit, S., G. Colletaz, C. Hurlin, and C. Pérignon (2013). A Theoretical and Empirical Comparison of Systemic Risk Measures. Working Paper.

Benoit, S., J.-E. Colliard, C. Hurlin, and C. Pérignon (2015). Where the risks lie: A survey on systemic risk. Working Paper.

[31] We could also calculate l_1 and use it as the initial leverage vector for the next round. However, drawing l_0 in step 4 independently from the pre-produced sample speeds convergence up as it avoids the otherwise strong autocorrelation of volatilities.

[32] Memmel and Raupach (2010) perform univariate estimates of capital ratios. They report a median monthly mean reversions of $\lambda_{\text{monthly}} = 0.18$. This AR(1) parameter on a monthly basis transfers into an Ornstein-Uhlenbeck mean reversion of $\lambda = -12 \times \log(1 - \lambda_{\text{monthly}}) = 2.38$, where the time unit is one year.

Bisias, D., M. Flood, A. W. Lo, and S. Valavanis (2012). A survey of systemic risk analytics. *Annual Review of Financial Economics 4*(1), 255–296.

Boucher, C. M., P. S. Kouontchou, and B. B. Maillet (2013). The Co-CoVaR and some other fair systemic risk measures with model risk corrections. Working Paper.

Brownlees, C. T. and R. F. Engle (2012). Volatility, correlation and tails for systemic risk measurement. Working Paper.

Collin-Dufresne, P. and R. S. Goldstein (2001). Do credit spreads reflect stationary leverage ratios? *Journal of Finance 56*, 1929–1957.

Danielsson, J., K. James, M. Valenzuela, and I. Zer (2014). Model Risk of Risk Models. Finance and Economics Discussion Series 2014-34, Board of Governors of the Federal Reserve System (U.S.), Washington, DC.

Fong, T. P. W. and A. Y. Wong (2012). Gauging potential sovereign risk contagion in Europe. *Economics Letters 115*(3), 496–499.

Gauthier, C., A. Lehar, and M. Souissi (2012). Macroprudential capital requirements and systemic risk. *Journal of Financial Intermediation 21*, 594–618.

Guntay, L. and P. H. Kupiec (2014). Taking the risk out of systemic risk measurement. Working Paper.

Hüser, A.-C. (2015). Too interconnected to fail: A survey of the interbank networks literature. Working Paper.

Leland, H. E. (1994). Corporate debt value, bond covenants, and optimal capital structure. *Journal of Finance 49*(4), 1213–52.

Leland, H. E. and K. B. Toft (1996). Optimal capital structure, endogenous bankruptcy, and the term structure of credit spreads. *Journal of Finance 51*(3), 987–1019.

López-Espinosa, G., A. Moreno, A. Rubia, and L. Valderrama (2012). Short-term Wholesale Funding and Systemic Risk: A Global Covar Approach. *Journal of Banking and Finance 36*(12), 3150–3162.

Meine, C., H. Supper, and G. N. Weiß (2015). Is tail risk priced in credit default swap premia? *Review of Finance*. forthcoming.

Memmel, C. and P. Raupach (2010). How do banks adjust their capital ratios? *Journal of Financial Intermediation 19*, 509–528.

Nucera, F., B. Schwaab, S. J. Koopman, and A. Lucas (2015). The information in systemic risk rankings. Working Paper.

Pericoli, M. and M. Sbracia (2003). A Primer on Financial Contagion. *Journal of Economic Surveys 17*(4), 571–608.

Stiroh, K. J. (2006). A Portfolio View of Banking with Interest and Noninterest Activities. *Journal of Money, Credit and Banking 38*(5), 1351–1361.

Straetmans, S. T. M., W. F. C. Verschoor, and C. C. P. Wolff (2008). Extreme US Stock Market Fluctuations in the Wake of 9|11. *Journal of Applied Econometrics 23*(1), 17–42.