DEALING WITH SYSTEMIC RISK WHEN WE MEASURE IT BADLY

JON DANIELSSON

KEVIN R. JAMES

MARCELA VALENZUELA

ILKNUR ZER

8 February 2012

Abstract: While an omniscient regulator would base a bank's capital requirement upon its contribution to systemic risk, we show that a regulator who measures a bank's contribution to systemic risk badly will find it optimal to use a simple leverage ratio instead. We empirically analyze the performance of leading risk measurement methods and find that they are incapable of providing either precise estimates of an individual bank's contribution to systemic risk or reliable rankings of banks by the amount of systemic risk they create. We conclude that a leverage ratio dominates a policy of systemic risk based capital requirements. JEL Codes: D81, G28, G32, G38.

* Corresponding author: Kevin James, Financial Markets Group, London School of Economics, Houghton St., London WC2A 2AE, UK. email: k.james1@lse.ac.uk. We thank the AXA Research Fund for financial support provided via the LSE Financial Market Group's program on risk management and financial regulation. We thank Bob DeYoung, Jean Helwege, Kostas Tsatsorinis and seminar participants at the Bank of England, the Bank of France, the Bank of Canada, the Financial Services Authority, the Reserve Bank of India, the Hong Kong Monetary Authority, the CCBS/LSE conference "Measuring Systemic Risk and Issues for Macroprudential Regulation" and the SUERF/Bundesbank Conference "The ESRB at 1" for helpful comments and discussions.

I. Introduction

Suppose that: i) a bank's stress or failure creates negative externalities for the economy as a whole by increasing the probability of a financial crisis (systemic risk); and ii) increasing the proportion of (costly) capital in the bank' s liability structure (that is, it's capital requirement) reduces this risk. It follows that an omniscient regulator would then require banks to finance themselves with capital to the point where the (observable to him) marginal cost of capital equals the marginal social benefits the capital creates by reducing the probability of a crisis. In other words, an omniscient being would base a bank's capital requirement upon its contribution to systemic risk (Crockett 2000). However, we ordinary mortals differ from omniscient beings (in the context of this essay) in that we must use highly imperfect models and data to estimate rather than observe a bank' s contribution to systemic risk. Consequently, we measure a bank's contribution to systemic risk imprecisely. Given this reality, do systemic risk based capital requirements make sense for us?

We answer this question in two steps. First, we show theoretically that a simple leverage ratio dominates a policy of basing bank capital requirement upon estimates of their contribution to systemic risk (the SysRisk policy) when those estimates are sufficiently imprecise. Second, we examine the empirical performance of the three leading approaches to measuring a bank's contribution to systemic risk, viz., Value-at-Risk, Marginal Expected Shortfall (Acharya et. al. 2010), and CoVaR (Adrian and Brunnermeier 2008). We show in each case that the proposed systemic risk measuring method (a riskometer, in the language of Danielsson 2009) produces both a very imprecise measure of an individual bank's absolute contribution to systemic risk and a completely unreliable ranking of banks by the level of systemic risk they create. In short, it is simply not possible to know with any degree of precision just how much systemic risk an individual bank creates in either absolute or relative (to other banks) terms. We therefore conclude that, while systemic risk based capital requirements (or other regulatory interventions) may make sense for an omniscient being, they probably do not make sense for regulators operating with the very limited information actually available to them.

We analyze the relative merits of the leverage ratio and a SysRisk policy by considering the case of a regulator who must set capital requirements for a banking system consisting of both high risk and low risk banks. In the absence of any information on individual bank risk, the best the regulator can do is to require all banks to operate with a leverage ratio appropriate for high risk banks (the HighCap policy) or low risk banks (the LowCap policy).[1] If he chooses the HighCap policy, he reduces the probability of a financial crisis (no high risk banks operate with too little capital), but low risk banks must finance themselves with a higher level of costly capital than is efficient. The LowCap option leads to the opposite trade-off (less costly capital, higher crisis probability). The regulator's choice between the HighCap and the LowCap policies then turns on whether it is less costly in total to treat high risk banks as low risk or low risk banks as high risk.

Turning now to the SysRisk policy, we assume that the regulator is endowed with an imperfect but better than random riskometer that provides a measure of each bank's risk. By using his riskometer and setting a bank's capital requirement on the basis of its reading, a higher proportion of banks are assigned the correct level of capital than is the case under the leverage ratio. However, riskometer imprecision sometimes leads to both high risk banks being treated as

[1] See D'Hulster (2009) for background on the leverage ratio.

low risk and low risk banks being treated as high risk, whereas the optimal leverage ratio policy leads to only the cheaper error. Thus, while the SysRisk policy leads to fewer classification errors, these errors are on average more costly. It follows that a regulator will find the SysRisk policy optimal only if his riskometer is precise enough to compensate for the more costly errors he makes while using it–a riskometer that is merely "better than nothing" will not suffice.

With this result in mind, we empirically evaluate the performance of the Value-at-Risk (VaR), Marginal Expected Shortfall (MES), and CoVaR riskometers. We base our sample upon the large US financial firm sample used in Acharya et. al. (2010), and we examine the pre-crisis (2003–2006) and crisis (2007–2010) periods separately. We assess riskometer performance along two dimensions, Reading Imprecision and Rank Imprecision. A riskometer's Reading Imprecision is low if it can pin down the level of a bank's risk to within a narrow range. A riskometer's Rank Imprecision is low if it can reliably order a sample of banks by the level of systemic risk they create (as defined by that riskometer). We assess riskometer precision by using a stationary bootstrap (Politis and Romano 1991) to construct 10,000 trial return histories for each sample period. We gauge Reading Imprecision by measuring the extent to which a given bank's risk reading varies across trials (i.e., a 99% confidence interval), and we gauge Rank Imprecision by seeing if a given riskometer provides a consistent ordering of banks by risk across trials.

We find that all three riskometers offer extremely imprecise readings along both performance dimensions. To begin with Reading Imprecision, the median confidence interval produced by the VaR riskometer spans 50 percentiles of the empirical distribution of the VaR readings for the sample banks. The MES and CoVaR riskometers perform even worse. In the case of MES (CoVar), the median confidence interval for a risk reading spans fully 65 (81) percentiles of empirical distribution of the MES (CoVaR) readings for sample banks. Turning to Rank Imprecision, we find for all three riskometers that the Spearman rank order correlation coefficient between the orderings of banks by risk in two trials selected at random is not statistically significantly different from 0. In short, the riskometers we have provide a poor foundation upon which to construct a policy of systemic risk based capital requirements.

We organize the remainder of this paper as follows. We review of the literature in Section II. We analyze the choice between a leverage ratio and a SysRisk policy as a function of riskometer precision in Section III. We examine the design of the three riskometers we focus upon in Section IV, and we analyze the empirical performance of these riskometers in Section V. Conclusions follow in Section VI.

II. Review Of The Literature

Crockett (2000) launched the quest for a systemic riskometer by arguing that a regulator should set a bank' s failure probability (via, for example, capital requirements) on the basis of that bank' s contribution to systemic risk. A measure of a bank's contribution to systemic risk would ideally be derived from a model of banks: i) in a complex economy; ii) that fully optimize; and iii) that pursue intricate state dependent strategies to do so. Elsinger, Lehar, and Summer (2002) pioneered the use of very detailed data on bank exposures to simulate the impact of shocks in a real (complex) economy, but to make their simulations tractable they had to assume that banks followed simple behavioral rules rather than optimal strategies (let alone complex optimal strategies). Goodhart, Sunirand, and Tsomocos (2004) argue that abrupt and non-linear changes in bank behavior in response to shocks play a crucial role in financial crises, and hence that one needs a general equilibrium model with optimizing banks to understand why and under what circumstances bank behavior will change in this manner. But, to make their analysis tractable, they had to work with highly stylized banks in a highly stylized economy. Danielsson, Shin, and Zigrand (2011) demonstrate theoretically that complex strategic interactions between financial market participants (endogenous risk) may play a crucial role in turning an exogenous shock into a financial crisis. However, their analysis is aimed at exploring the properties of a stylized financial system rather than at deriving a systemic risk measure for actual individual banks. So, while all of these approaches provide important qualitative insights into the nature of systemic risk, the bottom line for our purposes here is that there is now no model of complex optimizing banks in a complex economy that is capable of yielding a theoretically sound and empirically implementable measure of a bank's contribution to systemic risk.

In the absence of such a model, researchers devising practicable riskometers generally take the shock process (the initial shock plus any endogenous risk driven amplification/attenuation of that shock) as given and work with observable data on bank and financial system performance– with the idea being that performance data incorporate the impact of shocks that occur. While systemic risk measurement along these lines is a very active area of research (e.g., Segoviano and Goodhart 2009; Tarashev, Borio, and Tsatsaronis 2010; Zhou 2010), we focus here upon VaR, SES, and CoVaR, which are in our view the three most developed riskometers both theoretically and empirically.

Each riskometer we consider comes in an unconditional and a conditional (forecast) form. For example, Adrian and Brunnermeier (2008) develop both a conditional and an unconditional form of their CoVaR riskometer, and Brownlees and Engle (2011) develop the conditional version of MES. We choose to examine VaR, MES, and CoVaR in their unconditional forms. Our reason for doing so is that regulatory decisions on capital (etc.) necessarily evolve slowly given both the inherently time-consuming nature of regulatory decision-making and the time dependent adjustment costs to regulatory decisions on the part of regulated banks. Consequently, any decision involving capital (or other such regulatory variable) must be based upon a some sort of overall measure of a bank's risk rather than a measure of risk that depends a great deal upon the exact path of recent returns. Unconditional riskometers provide such an overall measure of risk, while the readings provided by conditional riskometers vary at far too high a frequency to be used as a basis for a regulatory requirement (though they may be useful in other contexts).

Our analysis of riskometer imprecision builds upon ideas developed by Green and Figlewski (1999) in their examination of the model risk facing a bank writing options and upon Danielsson's (2002; 2008) critiques of VaR. Danielsson (2002) explores the statistical assumptions behind the VaR measure, and Danielsson (2008) shows that the VaR riskometer performs poorly even in the relatively straightforward case of a $1000 investment in IBM. And while few would disagree with the observation that riskometers will inevitably be imprecise to at least some extent, we show here that the imprecision of one's riskometer matters as it affects the optimal policy choice

Hildebrand (2008) argues as we do that the uncertainties in risk measurement create a role for the leverage ratio. Our analysis differs from Hildebrand' s in the following respect: Hildebrand' s primary concern is that flaws in risk models and/or gaming of the capital rules by banks may leave banks with too little capital. Hence, he sees a role for the leverage ratio as a backstop for a SysRisk policy that ensures that all banks have some minimum level of capital. Admati et. al. (2011) extend Hildebrand's idea by arguing (on the basis of the Modigliani-Miller theorem) that the HighCap policy imposes little to no cost upon the banks while reducing the probability of a crisis, and hence that the HighCap policy is optimal. Our view is that both a SysRisk policy and a leverage ratio policy will lead to banks operating with inefficient levels of capital. The key to choosing between the two approaches then turns on the number and cost of the errors that each policy produces.

III. Optimal Policy When Using An Imprecise Riskometer

III.A. Set - Up and Assumptions

We examine the importance of riskometer precision by analyzing the case of a regulator who must set capital requirements for a banking system consisting of high risk and low risk (large) banks.[2] The trade-off the regulator faces is that while requiring banks to operate with a higher level of capital reduces the probability of a crisis, requiring banks to finance a higher proportion of their activities with capital is socially costly. The regulator must then choose a policy to set bank capital levels so as to minimize the total of the cost of the capital he requires banks to hold and the expected cost of a financial crisis.

We model this problem as a one period game in which the regulator first sets each bank's capital requirement and then a crisis occurs with a probability that is a function of the capital

[2] We assume that small banks do not create any systemic risk externalities and so do not include them in the analysis. In the context of the model, small banks can then operate with low levels of capital.

requirements he sets. We impose the constraint (consistent with actual behavior) that the regulator cannot act in a capricious and arbitrary fashion. We operationalize this constraint by assuming that the regulator must treat every bank that looks identical to him in an identical manner and must require all banks that look identical to him to hold the same level of capital.

We assume that the banking system consists of Q_H (Q_L) high risk (low risk) banks and to simplify the exposition we further assume that the number of banks is large enough to be treated as a continuum. All banks find it privately optimal to operate with a level of capital equal to K_L, and with this level of capital a low risk bank will never cause a crisis. However, if a high risk bank operates with a level of capital of less than K_H, $K_H > K_L$, it sets in train a series of events that ultimately causes a crisis with probability ϕ. Denote the number of high risk banks that operate with K_L as Q_{HL}. We assume that the cost of requiring a bank to operate with K_L is 0, that the cost of requiring a bank to increase its capital to K_H is a constant that does not depend upon the bank's type, and that the cost of requiring all banks to operate with K_H is κ.[3]

The social cost of a financial crisis is Ω, and the probability of a crisis is $\chi[Q_{HL}]$,[4] with

$$\chi[Q_{HL}] = 1 - (1-\phi)^{Q_{HL}} \quad (1)$$

[3] Alternatively, we could modify the model to enable banks to choose their level of risk as follows: high risk banks choose high risk behavior when operating with either K_H or K_L, and low risk banks choose low risk behavior when operating with K_L and high risk behavior when operating with K_H (at some social cost due to inefficiency). Changing the model in this manner will affect the optimal policy choice on the margin but it will not alter our principal result that the regulator prefers the SysRisk policy to the HighCap or LowCap policy only if his riskometer is sufficiently precise.

[4] We use the notation conventions of the computer program *Mathematica*, in which most of the work in the paper was done. So, $\chi[Q_{HL}]$ means that χ is a function of Q_{HL}.

It follows that $\partial\chi/\partial Q_{HL} > 0$ and that $\partial^2\chi/\partial Q_{HL}^2 < 0$. We assume that the social cost of a financial crisis is large and hence that a regulator would find it worthwhile to require all high risk banks to operate with capital equal to K_H if he could identify them. That is, we assume that

$$\Omega\chi\left[Q_H\right] > \kappa\frac{Q_H}{Q_H + Q_L} \qquad . \tag{2}$$

The regulator knows all of the parameters of the model, but he does not know which banks are high risk and which banks are low risk. The only way that he can get information on a bank's risk is to use his endowed riskometer. The riskometer will give one reading of either "High Risk" or "Low Risk" per bank, and that reading will be correct (incorrect) with probability π ($1-\pi$), with $0.5 < \pi \leq 1$.[5] That is, we assume that the riskometer is always better than random.

III.B. Optimal Policy

Since the requirement that the regulator treat all banks that look identical to him in an identical manner rules out any policy that involves a mixed strategy, the regulator has the following policy options available to him. He can eschew his riskometer (in which case all banks do look the same to him) and simply require that all banks operate with a uniform level of capital. The level can be set at either K_L (the LowCap policy) or at K_H (the HighCap policy). Alternatively, the regulator can choose to use his riskometer (the SysRisk policy), in which case he must use it for all banks and he must require a bank with a "High Risk" ("Low Risk") reading to operate with K_H (K_L). Consider the costs of each policy in turn.

---

[5] Intuitively, one could think that all banks try to appear to be low risk, and that the regulator can use his riskometer to (imperfectly) sort the actual high risk banks from the actual low risk banks. We have assumed for simplicity that the probability that the riskometer provides the correct reading is the same for both high risk and low risk banks, but allowing the probability to vary across bank types would not fundamentally alter the results below.

If the regulator chooses the LowCap policy, then all banks hold K_L at zero cost. It then follows that the number of high risk banks that hold K_L is just equal to the number of high risk banks, implying that the probability of a crisis is $\chi[Q_H]$. The total cost of this policy is then C_{LowCap}, with

(3) $$C_{LowCap} = \Omega\chi\left[Q_H\right].$$

If the regulator chooses the HighCap policy, then all banks hold K_H at a total cost of κ. The number of high risk banks that hold K_L is then zero, implying that the probability of a crisis is zero as well. The total cost of this policy is then $C_{HighCap}$, with

(4) $$C_{HighCap} = \kappa\ .$$

Note that neither $C_{HighCap}$ nor C_{LowCap} vary with π.

If the regulator chooses the SysRisk policy, then he treats each bank as if it is of the type indicated by his riskometer's reading. Consequently, the number of high risk banks operating with K_L is equal to the number of high risk banks that the riskometer mistakenly classifies as low risk, which is Q_H (1 - π). The total number of banks that must operate with K_H is equal to the number of high risk banks that the riskometer classifies correctly (Q_H π) plus the number of low risk banks the riskometer classifies incorrectly (Q_L (1 - π)). The total cost of the SysRisk policy is then $C_{SysRisk}$, with

(5) $$C_{SysRisk} = \Omega\chi\left[\left(1-\pi\right)Q_H\right] + \kappa\frac{Q_H\pi + Q_L\left(1-\pi\right)}{Q_H + Q_L}\ .$$

When π = 0.5, we have

(6) $$C_{SysRisk}\Big|_{\pi=0.5} = \Omega\chi\left[\frac{Q_H}{2}\right]+\frac{\kappa}{2},$$

and when π = 1, we have

(7) $$C_{SysRisk}\Big|_{\pi=1} = \kappa\frac{Q_H}{Q_H+Q_L}.$$

We sketch the total social costs of the regulator's policy options in Figure 1. Let us now consider the properties of and relationships between these cost functions.

LEMMA 1. $C_{SysRisk}\Big|_{\pi=0.5} > C_{SysRisk}\Big|_{\pi=1}$.

Proof. From equations 6 and 7 we know that $C_{SysRisk}\Big|_{\pi=0.5} > C_{SysRisk}\Big|_{\pi=1}$ holds if

(8) $$\Omega\chi\left[\frac{Q_H}{2}\right]>\frac{\kappa}{2}\frac{Q_H-Q_L}{Q_H+Q_L}.$$

Our assumption that the cost of a crisis is large (equation 2) and the relationship between the probability of a crisis and Q_H imply that

(9) $$\Omega\chi\left[\frac{Q_H}{2}\right]>\frac{\kappa}{2}\frac{Q_H}{Q_H+Q_L}.$$

It follows from equation 9 that the inequality in equation 8 holds.

∎

LEMMA 2. $\text{Min}\left[\left\{C_{LowCap}, C_{HighCap}\right\}\right] > C_{SysRisk}\Big|_{\pi=1}$.

Proof. We assume that the regulator will find it optimal to require high risk banks to operate with K_H if he can identify the high risk and low risk banks perfectly (equation 2), so it follows that $C_{LowCap} > C_{SysRisk}\big|_{\pi=1}$. We know that under both the HighCap policy and the SysRisk policy implemented with a perfect riskometer that all high risk banks operate with K_H, thereby eliminating the chance of a crisis. But, in order to achieve this benefit the HighCap policy requires both high risk and low risk banks to operate with K_H, while the SysRisk policy requires that only the high risk banks do so. Hence, $C_{HighCap} > C_{SysRisk}\big|_{\pi=1}$.

∎

LEMMA 3. $C_{SysRisk}\big|_{\pi=0.5} > \mathrm{Min}\left[\left\{C_{LowCap}, C_{HighCap}\right\}\right]$.

Proof. Suppose that the regulator prefers the HighCap policy to the LowCap. In this case, $C_{LowCap} \geq C_{HighCap}$, and it then follows from equations 3 and 4 that $\Omega\,\chi[Q_H] \geq \kappa$ and hence (with equation 1) that $\Omega\,\chi[Q_H/2] > \kappa/2$. Now compare the HighCap policy to SysRisk policy. It follows from equations 4 and 6 that $C_{SysRisk}\big|_{\pi=0.5} > C_{HighCap}$ if

$$\Omega\chi\left[\frac{Q_H}{2}\right] > \frac{\kappa}{2}, \tag{10}$$

which is true if $C_{LowCap} \geq C_{HighCap}$.

Now suppose that the regulator prefers the LowCap policy to the HighCap. In this case, $C_{HighCap} > C_{LowCap}$, and it then follows from equations 3 and 4 that

$$\frac{\kappa}{2} > \Omega\chi\left[\frac{Q_H}{2}\right], \tag{11}$$

Comparing the LowCap and the SysRisk policies, we find in this case that $C_{SysRisk}\big|_{\pi=0.5} > C_{HighCap}$

if

(12) $$\frac{\kappa}{2} > \Omega\chi\left[Q_H\right] - \Omega\chi\left[\frac{Q_H}{2}\right],$$

If follows from equation 1 that $\Omega\chi\left[\frac{Q_H}{2}\right] > \Omega\chi\left[Q_H\right] - \Omega\chi\left[\frac{Q_H}{2}\right]$ and then from equation 12 that

$C_{SysRisk}\big|_{\pi=0.5} > C_{HighCap}$.

∎

LEMMA 4. Consider a riskometer that returns the correct risk-reading with probability π**. If

$C_{SysRisk}\big|_{\pi=\pi^{**}} < C_{SysRisk}\big|_{\pi=0.5}$, then $C_{SysRisk}\big|_{\pi<\pi^{**}} > C_{SysRisk}\big|_{\pi=\pi^{**}}$ and

$C_{SysRisk}\big|_{\pi>\pi^{**}} < C_{SysRisk}\big|_{\pi=\pi^{**}}$.

Proof. Taking the derivative of $C_{SysRisk}$ with respect to π, we find that $\frac{\partial C_{SysRisk}\big|_{\pi=0.5}}{\partial \pi} \gtrless 0$ and that $\frac{\partial^2 C_{SysRisk}}{\partial \pi^2} < 0$. Suppose that $\frac{\partial C_{SysRisk}\big|_{\pi=0.5}}{\partial \pi} \leq 0$. In this case our result follows immediately. Now suppose that $\frac{\partial C_{SysRisk}\big|_{\pi=0.5}}{\partial \pi} > 0$. In this case it follows from Lemma 1 that: i) $C_{SysRisk}$ hits its maximum at π_M; ii) $\frac{\partial C_{SysRisk}\big|_{\pi>\pi_M}}{\partial \pi} < 0$; iii) there exists a π^* such that $C_{SysRisk}\big|_{\pi=\pi^*} > C_{SysRisk}\big|_{\pi=0.5}$; iv) $C_{SysRisk}\big|_{\pi<\pi^*} \geq C_{SysRisk}\big|_{\pi=0.5}$; and v) $\pi_M < \pi^* < 1$. Now consider $C_{SysRisk}\big|_{\pi^{**}>\pi^*}$. If $\pi > \pi^*$, we know from (ii) and (v) that $C_{SysRisk}\big|_{\pi>\pi^{**}} < C_{SysRisk}\big|_{\pi=\pi^{**}}$. If $\pi^* \leq \pi \leq \pi^{**}$, we know from (ii) and (v) that $C_{SysRisk}\big|_{\pi^*\leq\pi<\pi^{**}} > C_{SysRisk}\big|_{\pi=\pi^{**}}$. And, if $\pi < \pi^*$, we know from (iv) that $C_{SysRisk}\big|_{\pi<\pi^*} < C_{SysRisk}\big|_{\pi=\pi^{**}}$.

∎

Having established the key properties of the regulatory policy cost functions, we can now derive our main result.

PROPOSITION 1. The regulator prefers the leverage ratio to the SysRisk policy if his riskometer is sufficiently imprecise.

Proof. We know from Lemmas 2 and 3 that

$$C_{SysRisk}\big|_{\pi=0.5} > \text{Min}\left[\left\{C_{LowCap}, C_{HighCap}\right\}\right] > C_{SysRisk}\big|_{\pi=1}. \tag{13}$$

It follows that there exists a π^{**} such that $C_{SysRisk}\big|_{\pi=\pi^{**}} = \text{Min}\left[\left\{C_{LowCap}, C_{HighCap}\right\}\right]$. Lemma 4 then implies that the regulator prefers either the HighCap or the LowCap policy if $\pi < \pi^{**}$ and the SysRisk policy if $\pi > \pi^{**}$.

∎

Our analysis demonstrates that the regulator's optimal policy choice depends upon the quality of his riskometer. If the regulator finds himself in possession of a perfect riskometer, then a policy of setting capital requirements on the basis of that riskometer's readings is indeed optimal. However, as the regulator's riskometer becomes less precise, there comes a point where (poorly measured) systemic risk based capital requirements become dominated by other policies. So, what riskometers do we have and just how precise are they?

IV. RISKOMETER DESIGN: VaR, MES, And COVAR

A riskometer $\Psi\gamma$ is an algorithm that implements an approach γ for measuring a bank's contribution to systemic risk. Provided with the data concerning a given bank B that the riskometer requires, it provides a quantitative reading of B's contribution to systemic risk, denoted by $\Psi\gamma_{,B}$.

IV.A The VaR Riskometer

If one assumes that the marginal increase in the probability of a financial crisis due to a bank entering into a stressed state is independent of the state of the financial system, then a sensible measure of the systemic risk that a given bank creates is just the probability that the bank enters into a stressed state. The probability that a bank enters into a stressed state increases with the magnitude of a negative shock to its value (all else equal). Hence, banks that experience larger negative shocks create more systemic risk. If a bank does experience larger negative shocks, then one might expect its return distribution to have a longer left tail. A conceptually simple (if gameable) measure of the length of the tail is the position of a given extreme quantile of the return distribution. This line of thought leads to the VaR riskometer, where a bank's systemic risk measure is set equal to the value of a given extreme quantile of its return distribution.

To empirically implement the VaR riskometer, one must choose a return quantile and a method of estimating the location of that quantile. To begin with the quantile decision, the choice of exactly which quantile to use is of course arbitrary. The trade-off when selecting the quantile is that while a quantile further out in the tail provides a better estimate of the return consequences of a plausible worst case event, the further one goes out into the tail of the distribution, the fewer observations one has with which to estimate the location of the quantile for which one is aiming. The standard solution to this trade-off is to set the VaR quantile at either the 1^{th} percentile or the 5^{th} percentile of the return distribution. We use the 1^{th} percentile here.[6]

There exist a wide variety of methods that one can use to estimate the location of the VaR quantile one selects. To avoid needless complexity, we simply take as our estimate of the 1^{th} percentile of the return distribution the 1^{th} percentile value of the return series we examine (that is, we use the Historical Simulation method of estimating VaR).

We denote this quantile plus method specification of VaR as Ψ_{VaR}. So, Ψ_{VaR} for a bank B with a return series $\{R_{B,1}, R_{B,2} \ldots R_{B,T}\}$ is $\Psi_{VaR,B}$, with

$$(14) \qquad \Psi_{VaR,B} = \text{Quantile}\left[\left\{R_{B,1}, R_{B,1} \ldots R_{B,T}\right\}, 0.01\right].$$

We implement this measure using daily market returns for sample banks.

IV.B. The MES Riskometer

The MES riskometer takes its inspiration from the idea that, as Andrew Crocket (2000) put it, "the financial system is a system". That is, since the real economy requires the services that the financial system provides in order to function properly, a bank creates a risk to the overall economy (i.e., systemic risk) to the extent that it contributes to stress in the financial system as a whole. Building upon this idea, Acharya et. al. (2010) construct a theory yielding the result that

[6] We obtain the same results when using the 5^{th} percentile.

a bank creates systemic risk if it performs poorly at the same time as the economy as a whole performs poorly. This conclusion in turn leads to the MES riskometer, which is defined as a bank's expected return conditional upon the market performing poorly.[7] Acharya et. al. (2010) operationalize MES as follows: a bank B's MES is equal to its average return on days when the market return R_{Market} is at or below its 5th percentile value for the sample period ($VaR^{5\%}_{Market}$), determining the 5th percentile value of the market return series using the Historical Simulation method as above. Denoting this specification of MES as Ψ_{MES}, a bank B's Ψ_{MES} is then

$$\Psi_{MES,B} = \text{Mean}\left[\left\{R_{B,z_1}, R_{B,z_2} \ldots R_{B,z_W}\right\}\right] , \tag{15}$$

where $\{z_1, z_2 \ldots z_W\}$ is the set of days such that $R_{Market,z} < VaR^{5\%}_{Market}$. Following Acharya et. al. (2010), we compute this measure using daily returns and we use the return on the value-weighted CRSP index as our measure of R_{Market}.[8]

IV.C. The CoVaR Riskometer

CoVaR shares with MES the idea that a bank creates systemic stress when stress at the bank coincides with stress in the financial system as a whole, and it shares with VaR the idea that the location of the 1th percentile of a return distribution provides a good measure of risk. Combining these two ideas, Adrian and Brunnermeier (2008) posit that a bank B creates systemic risk if risk

---

[7] MES is the key building block of Acharya et. al.'s (2010) full systemic risk measure SES, which also includes a leverage measure. Given that SES builds upon MES, any imprecision in MES will necessarily affect SES. Thus, we concentrate our analysis here upon the MES part of SES.

[8] One aspect of MES that is worth noting is that, as Acharya et. al. (2010) acknowledge, the 5th percentile threshold is not really extreme enough to properly serve as a defining bound for "market stress". However, Acharya et. al. (2010) argue that this bound is far enough out into the tail of the market return distribution to lead to MES readings that are proportional to what MES readings would be if there was enough data to estimate them with a more realistically extreme bound to define episodes of market stress.

in the financial system as a whole increases as stress at the bank increases. Their measure of risk in the financial system is $VaR^{1\%}_{FinSys}$, which is the 1% VaR of the portfolio of financial firms.

Assume that B is under stress when $R_B = \Psi_{VaR,B}$ and not under stress when $R_B = \text{Median}[\boldsymbol{R}_B]$, where $\boldsymbol{R}_B$ is the return series for B. We then get the CoVaR riskometer's measure of systemic risk at bank B $\Psi_{CoVaR,B}$, with

$$\Psi_{CoVaR,B} = VaR^{1\%}_{FinSys}\Big|_{R_B=\Psi_{VaR,B}} - VaR^{1\%}_{FinSys}\Big|_{R_B=\text{Median}[\boldsymbol{R}_B]} \quad . \tag{16}$$

Following Adrian and Brunnermeier (2008), we estimate $VaR^{1\%}_{FinSys}\Big|_{R_B}$ with

$$VaR^{1\%}_{FinSys}\Big|_{R_B} = \alpha + \beta R_B + \varepsilon \quad . \tag{17}$$

using a quantile regression.[9] We then calculate $\Psi_{CoVaR,B}$ with the estimated parameters from equation 17.

There are obviously many ways of implementing this general CoVaR idea, and Adrian and Brunnermeier (2008) discuss several specifications in their paper. In particular, one must decide whether to measure returns by common stock returns or by changes in the market value of bank assets. Given that Adrian and Brunnermeier (2008) report that both methods lead to similar results in practice, we choose to measure returns with (easily observable) stock returns rather than by (difficult to infer) changes in the market value of bank assets so as to avoid a very

[9] We thank Johannes Ludsteck for sharing and updating the code he developed to estimate quantile regressions in *Mathematica*.

significant errors in variable problem.[10] While Adrian and Brunnermeier estimate Ψ_{CoVaR} with weekly returns, we use daily returns in our analysis below to be consistent with our analysis of Ψ_{VaR} and Ψ_{MES}.[11] We set R_{FinSys} equal to the return on the Fama-French Finance Industry portfolio.

V. RISKOMETER PERFORMANCE

V.A. Gauging Performance: Criteria and Method

Since the idea of the SysRisk policy is to make a bank's capital requirement a function of the amount of systemic risk that the bank creates, a regulator pursuing a SysRisk policy needs a riskometer that enables him to measure the amount of systemic risk that each individual bank creates. Ideally, a riskometer would produce a precise quantitative reading of each bank's risk, but at the very least the riskometer must be capable of ordering banks by the amount of risk they create. We therefore evaluate riskometer performance along two dimensions: Reading Imprecision and Rank Imprecision. A riskometer performs well on the Reading Imprecision dimension if the 99% confidence intervals for bank risk readings are "small". A riskometer performs well on the Rank Imprecision dimension if it provides a precise ordering of banks by the amount of systemic risk they create.

---

[10] The errors in variable problem can arise through several channels. Among these channels are the following. First, banks report asset values on a quarterly basis. One is therefore left with either very few return observations or one must use some ad hoc interpolation rule to trace the path of asset values over the quarter. Second, a considerable proportion of assets are reported at book rather than market value, so one must use some ad hoc method to transform book values to market values. Third, reported asset values are highly vulnerable to accounting manipulation. Fourth, publicly reported asset values do not include (at times substantial) off-balance sheet items and so may present a very incomplete picture of true assets.

[11] We repeated the analysis below using weekly returns and found that doing so did not alter our results.

Considering the Reading Imprecision dimension in more detail, our performance criteria requires that we have some way to judge whether the risk reading 99% confidence intervals produced by a given riskometer are "large" or "small". We propose to do so as follows. We take the meaning of "small" to be "small relative to the relevant yardstick", and we take as the relevant yardstick the distribution of the point estimate risk readings the riskometer in question produces for sample banks. That is, a riskometer's Reading Imprecision is low if it can pin down a bank's risk reading to within a narrow range of the distribution of point estimate risk readings.

To avoid the possibility of having this measure driven by extreme outliers and to make the measure more intuitively understandable, we report the range of a bank's risk reading confidence interval in terms of percentiles of the point estimate distribution. To illustrate, suppose that the point estimate risk readings that a given riskometer produces for a sample of banks are distributed according to a Uniform distribution with endpoints {–5, –10} and that the 99% confidence interval for bank *B*'s risk reading is {–6, –9}. In this case, *B*'s Reading Imprecision Bounds are {20, 80} and its Reading Imprecision Score is 60 (i.e., 80 - 20) as its confidence interval stretches from the 20th percentile to the 80th percentile of the distribution of the point estimate risk readings. The maximum value of a Reading Imprecision Score is 100 (meaning that the bank's risk reading could be anywhere in the point estimate distribution) and the minimum value is 0 (meaning that the reading is perfectly precise).

As there is no simple analytical formula available to compute confidence intervals for readings produced by the riskometers we examine, we calculate confidence intervals using a bootstrap. For each sample period, we generate 10,000 return histories ("trials"). These trials produce 10,000 readings for each bank/riskometer combination. To get the 99% confidence interval for a given bank's risk reading for a given riskometer, we sort that bank's riskometer readings by value and set the confidence interval equal to the bounds of the middle 99% of that distribution. To measure Rank Imprecision for a given riskometer, we calculate the Spearman rank order correlation coefficient ρ for 10,000 pairs of trials selected at random, where the order of each bank in each trial is determined by its reading for that riskometer. We then get the distribution of ρ for that riskometer, and we see if it is statistically significantly different from 0.

Since our return data incorporates both cross-sectional and time series relationships, our trials must capture these aspects of the data. We therefore generate the trials as follows. To capture the cross-sectional relationships between R_{Market}, R_{FinSys}, and individual bank returns, we construct our trials by drawing sample days. So, if bank B's return has a cross-sectional relationship with R_{Market}, that relationship will appear in the trials as we will have R_{Market} and R_B observations for the same days. To capture the time series relationships in the data, we select the days we include in a trial using a stationary bootstrap (Politis and Romano 1991). In this technique one initially draws an observation day at random from the sample and then includes a random number of following days. The number of additional days that one includes is given by a draw from a geometric distribution with parameter g, with $g = T^{-1/3}$, where T is the number of days in the sample period. One then repeats this process until the number of days in the trial equals T.

V.B. Data

Systemic risk concerns arise in connection with large financial firms. We therefore construct our sample by beginning with the large financial firm sample used in Acharya et. al. (2010). This sample consists of the 101 financial firms with a market cap in excess of $5 billion as of

June 2007. We examine a pre-crisis sample period consisting of the years 2003–2006, and a crisis sample period consisting of the years 2007–2010. We include a bank in a sample period if we have a return observation for that bank for at least 75% of sample days. Our 2003–2006 sample consists of 92 firms, and our 2007–2010 sample consists of 77 firms. We obtain our firm returns from CRSP.

V.C. Reading Imprecision: VaR

Carrying out the analysis discussed above, we plot Ψ_{VaR} 99% confidence intervals and Reading Imprecision Bounds for each sample period in Figure II and we report summary statistics on the Reading Imprecision Scores in Table I. As one would expect, the absolute magnitude of Ψ_{VaR} confidence intervals soared in the crisis period. But, since the Ψ_{VaR} point estimate distribution stretched by a comparable amount, the Reading Imprecision Bounds remained roughly constant. That is, Ψ_{VaR} is as good–or as bad–a riskometer in the crisis period as it is in the pre-crisis period.

That said, the Reading Imprecision Bounds and Scores indicate that Ψ_{VaR} does not perform all that well. The median Reading Imprecision Score is 51 (52) in the 2003–2006 (2007–2010) sample, indicating that the 99% confidence interval of the typical bank's Ψ_{VaR} reading spans fully half of the Ψ_{VaR} distribution. Equipped with the Ψ_{VaR} riskometer, then, a regulator would find it difficult to correctly assign the typical bank to even the correct quartile of the risk distribution.

V.D. Reading Imprecision: MES

Turning now to the more theoretically sound but more complex MES riskometer, we plot Ψ_{MES} 99% confidence intervals and Reading Imprecision Bounds in Figure III, and we report summary statistics on Reading Imprecision Scores in Table I. In the case of Ψ_{MES}, too, the risk reading confidence intervals increase dramatically in the crisis period, but the Reading Imprecision Bounds and Scores remain roughly constant. Again, though, our Reading Imprecision results indicate that Ψ_{MES} performs poorly as a riskometer. In this case the median Reading Imprecision Score is 65 (68) for the 2003–2006 (2007–2010) sample, indicating that the 99% confidence interval for the MES riskometer's reading for the typical bank spans approximately two thirds of the Ψ_{MES} point estimate distribution. A regulator equipped with the MES riskometer would often find it impossible to reliably tell if a given bank belonged in either the upper or the lower quartile of the risk distribution.

V.E. Reading Imprecision: CoVaR

CoVaR shares with MES an elegant theoretical foundation, but Ψ_{CoVaR} is even more complex to estimate. The CoVaR riskometer's reading for a given bank is produced by a combination of two items that are each very difficult to measure on their own, viz., the Ψ_{VaR} of the bank and $VaR^{1\%}_{FinSys}$ as a function of the bank's return. Given that our results above demonstrate that it is impossible to measure even Ψ_{VaR} alone with precision, it will come as no surprise to see that Ψ_{CoVaR} is the least precise of the riskometers we examine.

We plot Ψ_{CoVaR} 99% confidence intervals and Reading Imprecision Bounds in Figure IV, and report Ψ_{CoVaR} Reading Imprecision Scores in Table I. The plots indicate and the Imprecision Scores confirm that the CoVaR riskometer is so imprecise that a reading conveys very little information about the location of a bank in the risk distribution. The median Reading Imprecision Score in the 2003–2006 (2007–2010) is 82 (95). Thus, a regulator equipped with the CoVaR riskometer could not reliably tell if the typical bank belonged in either the upper or the lower decile of the risk distribution.

V.F. Rank Precision

Our Reading Imprecision analysis demonstrates that individual bank riskometer readings vary tremendously across trials. But, it could be the case that individual bank riskometer readings covary from trial to trial as well. That is, some trials may produce high riskometer readings for all banks, and some trials may produce low riskometer readings for all banks. So, our results on Reading Imprecision above do not rule out the possibility that a regulator can rank banks by the amount of risk they create. If this is the case, then a regulator would at least find it possible to require high risk banks to hold more capital than low risk banks.

We examine this possibility by computing the Spearman rank order correlation coefficient ρ between trials selected at random. If banks tend to appear in the same order when ranked on the basis of riskometer readings from trial to trial (even if the value of those readings varies a great deal), then ρ will be positive. We obtain the distribution of ρ by computing it for 10,000 pairs of trials, with each trail in the pair selected at random. We sort the set of ρ's we obtain by value and take the 99% confidence interval to be equal to the bounds on the middle 99% of that distribution.

We report the results of this analysis in Table II. We find for each riskometer and each sample period that the ranking of banks by riskometer reading are essentially uncorrelated from trial to trial, and hence that ρ is not statistically significantly different from 0. It follows that a regulator can not use the riskometers we examine here to order banks by the level of systemic risk they create.

VI. CONCLUSION

Basing a bank's capital requirement (or other regulatory intervention) upon the level of systemic risk that the bank creates is undoubtedly far superior to basing capital requirements upon a cruder policy such as a simple leverage ratio...at least in theory. Yet, the SysRisk policy will only be superior in practice if the regulator possesses a riskometer that enables him to precisely measure the amount of systemic risk that individual banks create. We examine the three leading riskometers, viz., VaR, MES, and CoVaR, and find that they are incapable meeting the demands that a SysRisk policy places upon them. We therefore conclude that, in practice, less informationally intensive policies such as the leverage ratio may offer a more sensible approach to dealing with systemic risk than a more theoretically attractive SysRisk policy.

We note that our critique of the riskometers we examine is not that they are theoretically unsound or that they are mis-specified. If this were the case, then a different riskometer or a different specification for an existing riskometer might solve the problem. But this is not the case: MES and CoVaR in particular have strong theoretical foundations, and these foundations lead to plausible empirical risk measures. Instead, we show that the three analytically distinct riskometers that we do examine are all empirically unreliable. Since any riskometer will have to confront similar estimation issues, we think it likely that our conclusion that it would be difficult to build a SysRisk policy upon the riskometers we examine here extends to riskometers more generally.

We draw two implications from our analysis. First, on a methodological note, our analysis highlights the importance of taking measurement uncertainty seriously. A policy option that may make sense for an omniscient being that can observe all of the relevant parameters directly may not make sense for regulators who must operate with imperfect estimates of those parameters. Second, while we do think that the leverage ratio may dominate the SysRisk policy in practice, this is not because the leverage ratio is itself ideal–it is just that we don't know enough to get the SysRisk policy to work well. It follows that capital will inevitably be an inefficient tool with which to address systemic risk. And while the inefficiency of capital by no means implies that a better tool exists, it does suggest that it would be worth looking hard to find one (or more).

Our analysis of systemic risk based capital requirements brings to mind the story of Phaeton, the long lost son of the god Apollo. One day Phaeton appeared on Mount Olympus and Apollo was so delighted to see him that he offered to grant him anything that he desired. Phaeton asked to drive the chariot that pulled the Sun across the heavens. Apollo argued that this was a terrible idea as the chariot was built for him–a god–and not for a mortal. But Phaeton stood firm and Apollo had given his word, so the next morning Phaeton set off. Of course, he was completely incapable of controlling Apollo's chariot and was thus in the process of destroying the world when Zeus stuck him down with a thunderbolt. The moral of the story is that one should be wary of using divine tools if one lacks divine powers. This is wisdom that we forget at our peril.

LONDON SCHOOL OF ECONOMICS

LONDON SCHOOL OF ECONOMICS

LONDON SCHOOL OF ECONOMICS

REFERENCES

Acharya, Viral, Lasse Pederson, Thomas Philippon, and Matthew Richardson, "Measuring Systemic Risk", NYU Working Paper, 2010

Admati, Anat, Peter DeMarzo, Martin Hellwig and Paul Pfleiderer, "Fallacies, Irrelevant Facts, and Myths in the Discussion of Capital Regulation : Why Bank Equity is Not Expensive", Stanford GSB Working Paper, 2011

Adrian, Tobias and Marcus Brunnermeier, "CoVaR", Federal Reserve Bank of New York Working Paper no. 348, 2008

Brownlees, Christian and Robert Engles, "Volatility, Correlation, and Tails For Systemic Risk Measurement", NYU Working Paper, 2010

Brunnermeier, Markus, Andrew Crockett, Charles Goodhart, Avinash Persaud, and Hyun Shin, "The Fundamental Principles of Financial Regulation", *Geneva Reports on the World Economy,* 2009

Crocket, Andrew, "Marrying the Micro - and the Macro - Prudential Dimensions of Financial Stability", BIS Speech, 2000

D' Hulster, Katia, "The Leverage Ratio", *World Bank Crisis Response,* Note 11, 2009

Danielsson, Jon, "The Emperor Has No Clothes : Limits to Risk Modelling", *Journal of Banking and Finance,* 26 (2002), 1273–1296

______, "Blame the Models", *Journal of Financial Stability,* 4 (2008), 321–328

______, "The myth of the riskometer", VoxEU.org, http://voxeu.org/index.php?q =node/2753, 2009

Danielsson, Jon, Hyun Shin, and J.P. Zigrand, "Endogenous and Systemic Risk", http://www.riskresearch.org, 2011

Efron, Bradley and R.J.Tibshirani (1994), *An Introduction to the Bootstrap*, 1st ed. (New York, NY: Chapman and Hall, Monographs on Statistics and Applied Probability 57, 1993)

Elsinger, Helmut, Alfred Lehar, and Martin Summer, "Risk Assessment for Banking Systems", Oesterreichsiche NationalBank Working Paper No.79, 2002

Goodhart, Charles, Pojanart Sunirand and Demitrios Tsocomos, "A Model to Analyse Financial Fragility : Applications", Journal of Financial Stability 1 (2004), 1 - 30

Green, T.Clifton and Stephen Figlewski (2009), "Market Risk and Model Risk for a Financial Institution Writing Options", *Journal of Finance,* 54 (2009), 1465 - 1499

Hildebrand, Philipp, "Is Basel II Enough?The Benefits of the Leverage Ratio", http://www.bis.org/review/r081216d.pdf, 2008

Politis, Dimitris N. and Joseph P. Romano (1991), "The Stationary Bootstrap", Technical Report #91–03, Department of Statistics, Purdue University, 1991

Segoviano, Miguel and Charles Goodhart, "Banking Stability Measures", *IMF Working Paper* WP/09/04, 2009

Tarashev, Nikola, Claudio Borio and Kostas Tsatsaronis, "Attributing Systemic Risk to Individual Institutions", BIS Working Paper 308, 2010

Zhou, C., "Are banks too big to fail? Measuring systemic importance of financial institutions", *International Journal of Central Banking* 6 (2010), 205–250

TABLE I

Reading Imprecision Scores

Sample	Riskometer	Mean	Median	Standard Deviation	Minimum	Maximum
2003/2006	VaR	49	51	17	7	90
	MES	57	65	21	3	93
	CoVaR	81	82	16	23	100
2007/2010	VaR	49	52	19	4	78
	MES	62	68	19	8	97
	CoVaR	93	95	7	43	100

Notes: In this table we report summary statistics on the Reading Imprecision Scores for the three riskometers we examine for each sample period. A bank's Reading Imprecision Score for a given riskometer is equal to the range of the 99% confidence interval of that riskometer's reading for that bank, with the confidence interval bounds expressed in terms of percentiles of the empirical distribution of the point estimate riskometer readings for sample banks. To illustrate, if the point estimate riskometer readings for the sample were uniformly distributed on {–5, –10} and the bounds of the 99% confidence interval for a given bank's risk reading were {–6, –9}, then that bank's Reading Imprecision Score would be 60 (80 – 20) as its confidence interval spans the 20th to the 80th percentiles of the distribution of sample bank riskometer readings. Our base sample consists of the 101 financial firms with a market cap in excess of $5 billion as of June 2007. We include a firm in a sample period if we have a return observation for that firm for at least 75% of sample days. Our 2003/2006 sample consists of 92 firms, and our 2007/2010 sample consists of 77 firms.

TABLE II

Rank Imprecision

Sample	Riskometer	Median	99% Confidence Interval
2003/2006	VaR	0.05	{-0.22, 0.32}
	MES	0.01	{-0.26, 0.28}
	CoVaR	0.03	{-0.24, 0.29}
2007/2010	VaR	0.12	{-0.19, 0.41}
	MES	0.09	{-0.21, 0.39}
	CoVaR	0.08	{-0.2, 0.36}

Notes: In this table we report summary statistics for the Spearman rank order correlation coefficient for each riskometer for each sample period. We compute these statistics by: i) using a stationary bootstrap to contruct 10,000 trial return histories for each sample; ii) ordering the sample firms by riskometer readings for each trial; and iii) calculating the Spearman coefficient for 10,000 pairs of trials selected at random. Our base sample consists of the 101 financial firms with a market cap in excess of $5 billion as of June 2007. We include a firm in a sample period if we have a return observation for that firm for at least 75% of sample days. Our 2003/2006 sample consists of 92 firms, and our 2007/2010 sample consists of 77 firms.

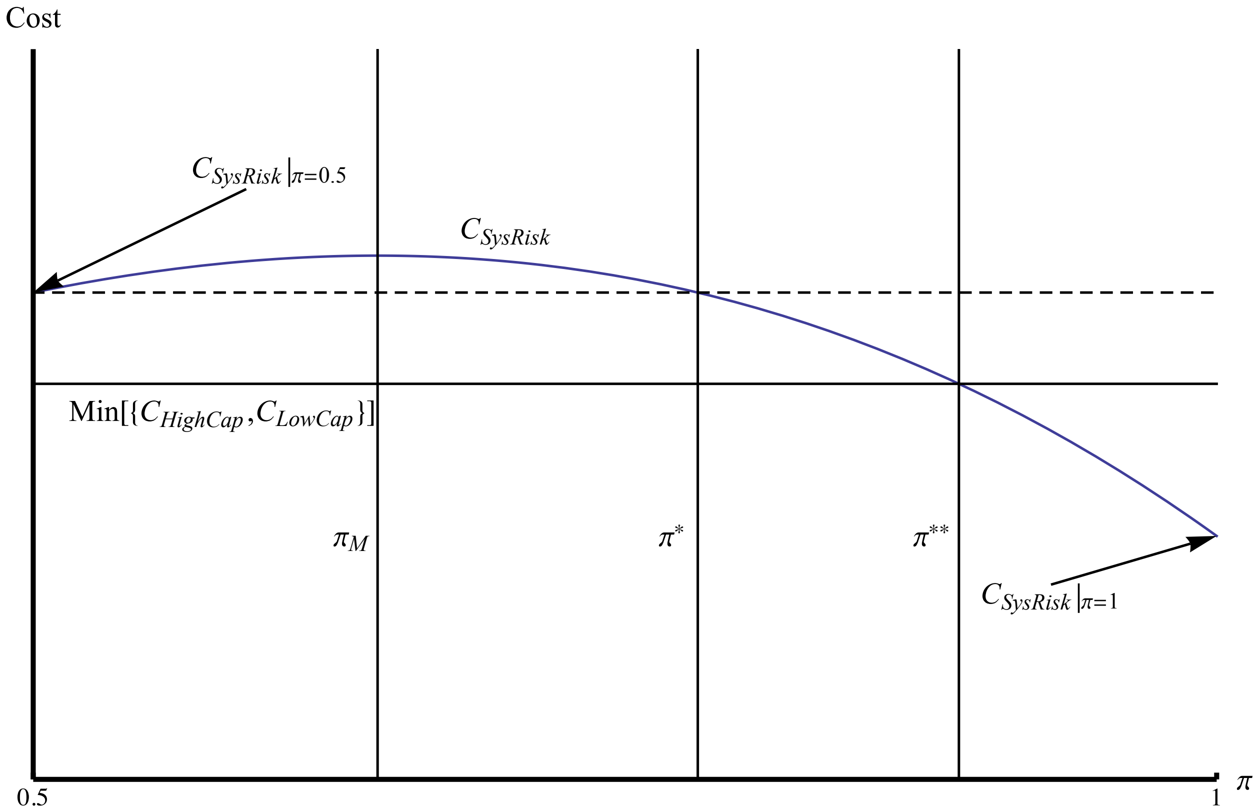

FIGURE 1

Riskometer Precision and the Optimal Policy Choice

In this figure we sketch the cost of the SysRisk policy (the $C_{SysRisk}$ curve) and the cost of the best of the two leverage ratio policies HighCap and LowCap (indicated by the $\text{Min}\left[\left\{C_{HighCap}, C_{LowCap}\right\}\right]$ line) as a function of the precision of the regulator's riskometer (π). π_M indicates the value of π where $C_{SysRisk}$ is at its maximum, π^* indicates the value of π such that $C_{SysRisk}\big|_{\pi=\pi^*} = C_{SysRisk}\big|_{\pi=0.5}$, and π^{**} indicates the value of π such that

$$C_{SysRisk}\big|_{\pi=\pi^{**}} = \text{Min}\left[\left\{C_{HighCap}, C_{LowCap}\right\}\right] .$$

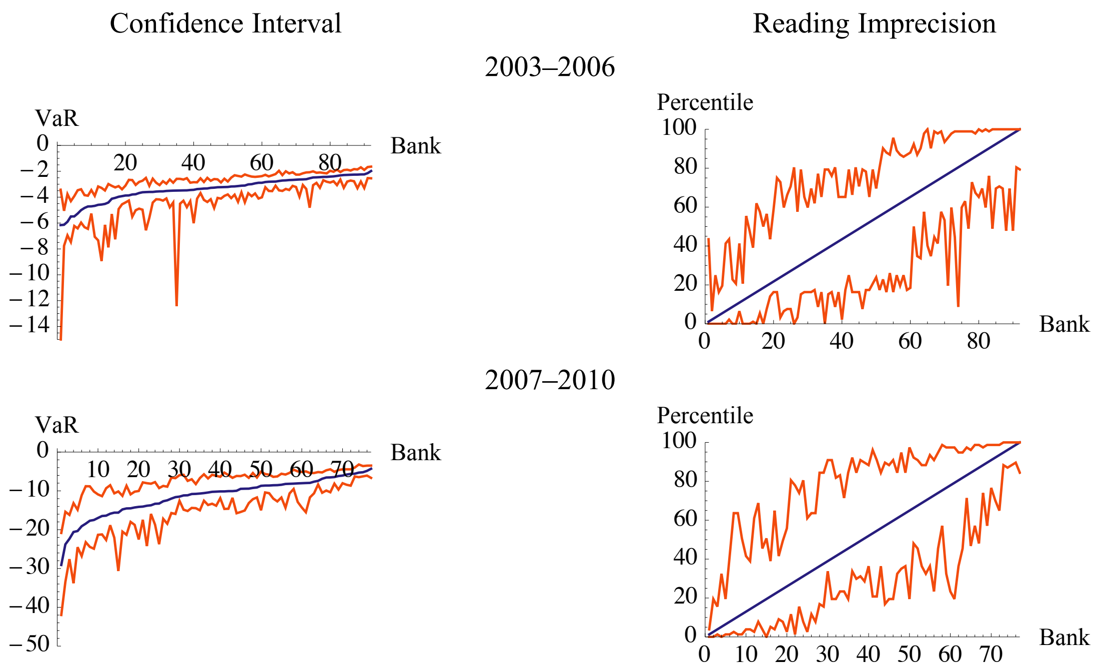

FIGURE II

VaR Reading Imprecision

In this figure we plot the point estimate and the 99% confidence interval of each sample bank's VaR estimate on the left and the implied Reading Imprecision bounds on the right for each sample period, with a bank's VaR equal to the 1^{th} percentile of its one day return distribution. The Reading Imprecision plot expresses the VaR point estimates and confidence intervals in terms of the percentile location of their values in the empirical distribution of the VaR point estimates.

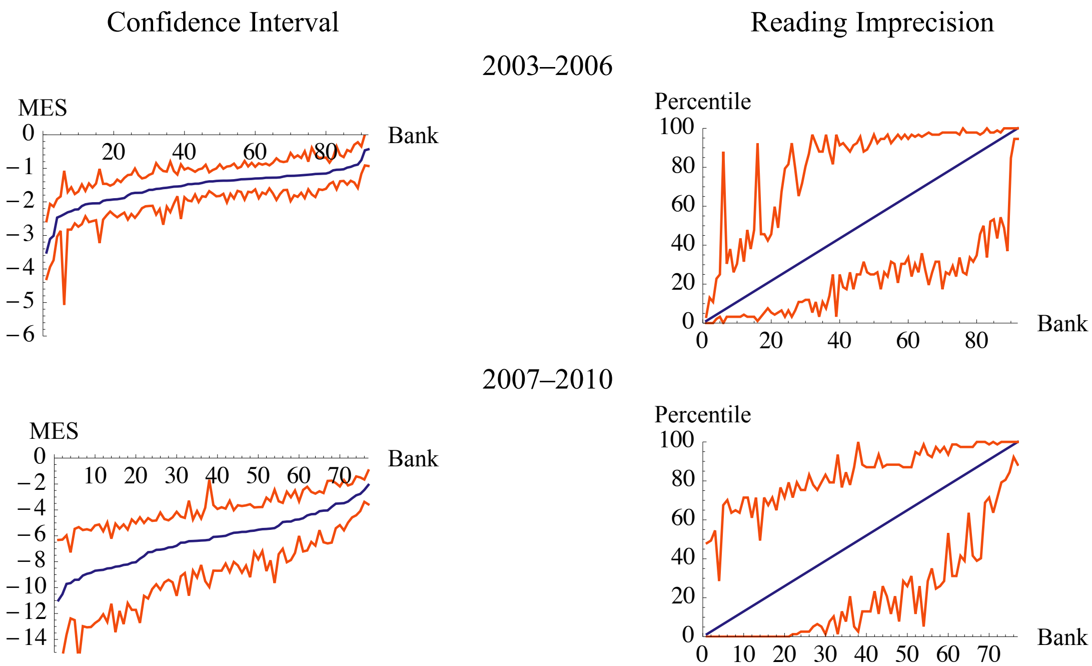

FIGURE III

MES Reading Imprecision

In this figure we plot the point estimate and the 99% confidence interval of each sample bank's MES estimate on the left and the implied Reading Imprecision bounds on the right for each sample period, with a bank's MES is equal to its average return on days when the return on the market portfolio is below its 5th percentile value for the sample period. The Reading Imprecision plot expresses the MES point estimates and confidence intervals in terms of the percentile location of their values in the empirical distribution of the MES point estimates.

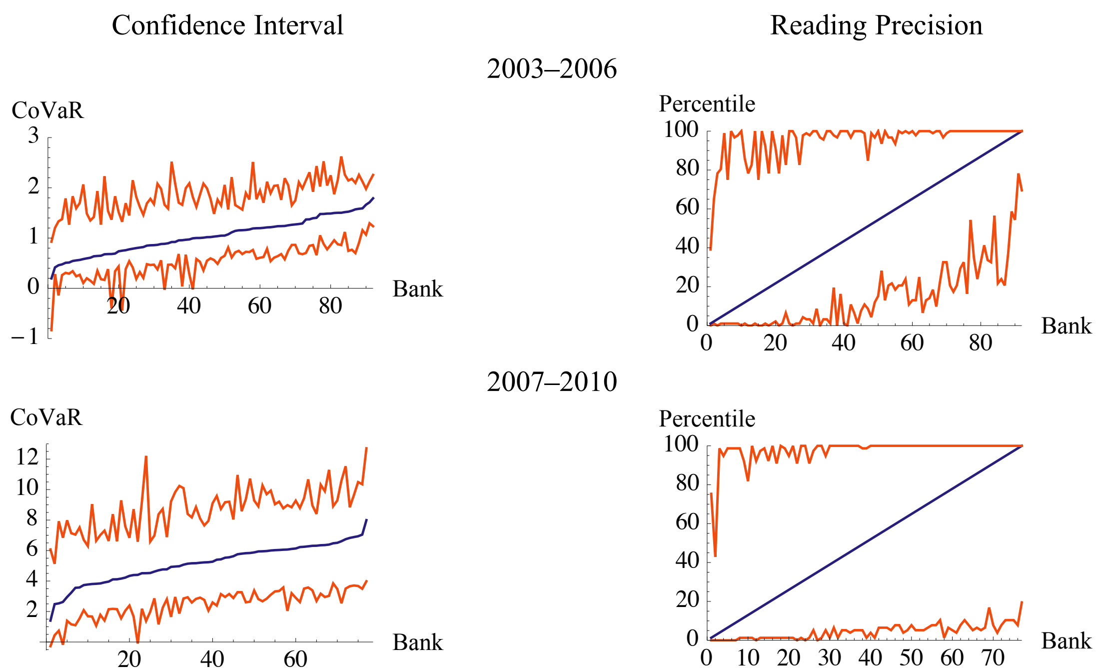

FIGURE IV

CoVaR Reading Imprecision

In this figure we plot the point estimate and the 99% confidence interval of each sample bank's CoVaR estimate on the left and the implied Reading Imprecision bounds on the right for each sample period. A bank's CoVaR is equal to the difference between the VaR of the portfolio of financial sector firms given that the return on the bank equals its median value and the VaR of the portfolio of financial sector firms given that the return on the bank equals its 1th percentile value, with the VaR of the portfolio of financial sector firms equal to the 1th percentile of its one day return distribution. The Reading Imprecision plot expresses the CoVaR point estimates and confidence intervals in terms of the percentile location of their values in the empirical distribution of the CoVaR point estimates.