Mathias Drehmann
mathias.drehmann@bis.org

Nikola Tarashev
nikola.tarashev@bis.org

Systemic importance: some simple indicators[1]

Are there simple yet reliable indicators of banks' systemic importance? In addressing this question, this article explores three model-based measures of systemic importance and finds that bank size helps approximate each of them. A bank's total interbank lending and borrowing provide useful complementary information.

JEL classification: G20, G28, L14.

A pressing policy objective is to finalise and implement a regulatory framework for systemically important financial institutions. Meeting this objective calls for measures of systemic importance. The recent academic literature has proposed a number of such measures, underpinned by sophisticated economic and statistical techniques. Despite their intellectual appeal, these measures pose serious challenges for practitioners. They are demanding on data, computationally intensive and difficult to communicate to the general public. In addition, given that the measures require detailed system-level information, individual institutions would not be able to use these measures directly in order to assess and manage their own degree of systemic importance. This prompts the question whether there are simple, readily available indicators that are reliable proxies for the more sophisticated measures.

In this article, we address this question empirically. We use data on 20 large internationally active banks to test the relationship between three sophisticated, model-based measures of systemic importance and three simple indicators. Given the multifaceted nature of systemic importance, we consider both top-down and bottom-up measures. The top-down measures first derive systemic (ie system-wide) risk and then allocate it to individual institutions. We explore two such measures that differ in terms of their perspective on systemic importance and, consequently, in the way in which they allocate system-wide risk. We also consider one bottom-up measure, which first assumes distress in a particular institution and then evaluates the level of system-wide risk associated with that event. We then compare each of these measures to simple indicators that are

[1] The views expressed are those of the authors and do not necessarily reflect those of the BIS. Jörg Urban and Marek Hlaváček provided excellent research assistance. Claudio Borio, Stephen Cecchetti, Robert McCauley, Kostas Tsatsaronis and Christian Upper provided helpful comments on earlier drafts of the article.

based on readily available and well understood characteristics of individual banks: size, total interbank lending and total interbank borrowing.

We find that the simple indicators approximate the model-based measures of systemic importance quite well. Under each of these measures, bank size is highly significant in both statistical and economic terms. In comparison, the link between interbank activity and measured systemic importance is weaker. And the strength of this link varies across the alternative measures of systemic importance, in line with the economic logic underlying each of them.

We perform the analysis in three steps. We start with a specific definition of system-wide risk, which is necessary for constructing any measure of systemic importance. We then outline three such measures, highlighting the different perspectives on systemic importance they incorporate. Finally, after describing our empirical setup, we evaluate the explanatory power and economic significance of the simple indicators for each of the three rigorous measures.

A measure of systemic risk

A popular measure of systemic risk …

Systemic risk is an elusive concept: it can have significant economic consequences and is quantitatively important, yet there is no clear consensus on how it should be measured. In this article, we associate systemic risk with losses in the financial system exceeding a high threshold with a small probability. We regard such losses as indicating systemic events, which are also characterised by a disruption to financial services and potentially serious harm to the real economy (FSB, IMF, BIS (2009)).

In measuring systemic risk, we only consider losses incurred by banks' *non-bank* creditors as opposed to bank equity holders and interbank creditors. Thus, our perspective is that of an insurer of banks' debt whose concern is solely about system-wide losses vis-à-vis the rest of the economy. By abstracting from losses to equity holders, we treat equity as fully loss-absorbing. In other words, a positive equity value, no matter how small, ensures the smooth functioning of a bank and does not imply any systemic repercussions.[2] In turn, by abstracting from losses to interbank creditors, we avoid double-counting. Since the interbank liabilities of one bank are the interbank assets of another, losses to the interbank creditors of one bank are losses incurred by the equity holders or non-bank creditors of one or more other banks in the system.

Concretely, we measure systemic risk as the expected *aggregate* loss to non-bank creditors, conditional on such a loss exceeding the 99th percentile of the underlying probability distribution. This measure is often referred to as the expected shortfall at the 99% confidence level.

[2] Admittedly this is a strong assumption, given that the recent financial crisis showed that banks' strategic behaviour at positive but low equity values can have adverse systemic consequences. The assumption also abstracts from informational frictions associated with losses on interbank positions or the possibility that banks have cross-shareholdings. That said, postulating that equity is a true loss absorber only above a certain level would increase the magnitude of measured losses but would not change the main messages of this article.

Expected shortfall (ES) is the most popular measure of systemic risk. Its popularity stems from the fact that, unlike most of its alternatives, it provides an informative summary of the severity of extreme events that occur with a small probability but can have system-wide consequences. Recent studies by Acharya et al (2009), Webber and Willison (2011) and Huang et al (2010) apply ES to the analysis of systemic risk in a variety of settings.

Measures of systemic importance

... underpins three rigorous measures of systemic importance

A bank's systemic importance can be measured from different angles. Like slicing a pie into pieces, top-down measures start with the risk of the system and allocate it to individual institutions. By contrast, bottom-up measures start with distress at a particular institution and then compute the associated level of system-wide distress. Even when based on the same measure of systemic risk, as is the case here, these alternative measures of banks' systemic importance typically deliver different conclusions. It is thus useful to explain the intuition behind the underlying approaches.

Top-down measures

The literature has proposed two approaches to allocate systemic risk across banks in a top-down fashion. The first, the participation approach (PA), considers the expected participation of each bank in systemic events. As illustrated by Graph 1, PA first focuses on systemic events (shaded area in the left-hand panel). It then measures the systemic importance of a bank, say bank *i*, as the expected losses incurred by its non-bank creditors in these events. Economically, PA equals the actuarially fair premium that the bank would have to pay to a provider of insurance against losses it may incur in a systemic event.[3]

As argued in Tarashev et al (2010) and Drehmann and Tarashev (2011), however, the extent to which a bank participates in systemic events typically differs from the extent to which it *contributes* to these events. To see why,

Participation approach (PA)

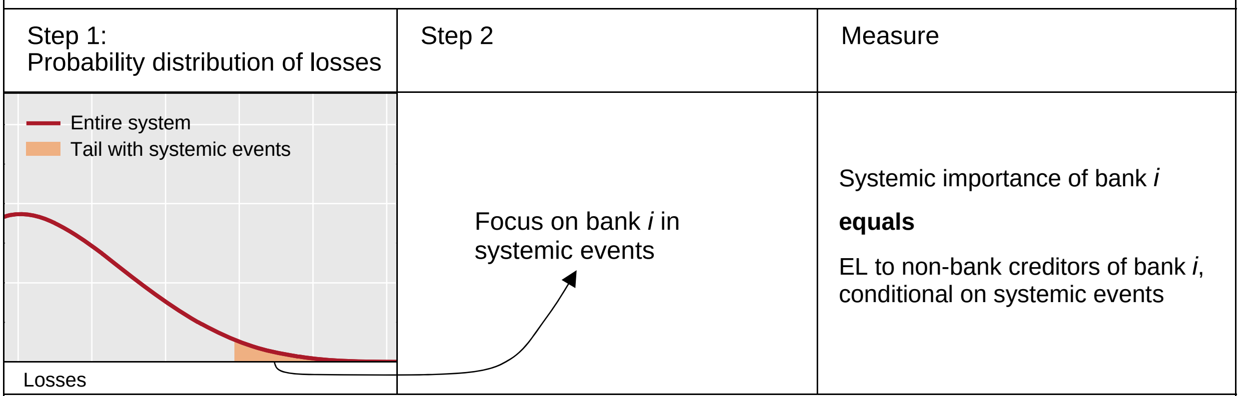

EL = expected loss.

Graph 1

[3] PA has been implemented in various ways and with different datasets by Acharya et al (2009), Huang et al (2010), Brownlees and Engle (2010), Tarashev et al (2010) and Drehmann and Tarashev (2011).

consider a bank with small debt liabilities to non-banks but with large positions on the interbank market. Since the failure of this bank in a systemic event would impose small losses on non-banks, we say that it participates little in systemic events. But the bank may contribute materially to these events by transmitting distress from one bank in the system to another. Such cases are captured by the second top-down approach: the contribution approach (CA). CA accounts explicitly for the fact that a bank contributes to systemic risk through its exposure to exogenous shocks, by propagating shocks through the system, and by being itself vulnerable to propagated shocks.[4]

CA is rooted in a methodology first proposed by Shapley (1953) for the allocation across individual players of the value created in a cooperative game. The intuition behind this methodology is quite simple. We could use the level of risk an individual bank generates in isolation as a measure of systemic importance. But such an approach would miss the contribution of each bank to the risk of others. Similarly, it is not enough to consider only the marginal-risk contribution of a single bank, calculated as the difference between the system-wide risk with and without the bank. The reason is that this calculation ignores the complexity of bilateral relationships, which is especially pronounced when interbank exposures can propagate shocks within the system through a potentially long chain of market participants. The Shapley methodology accounts fully for such interactions by ascribing to individual institutions a weighted average of the marginal contributions each makes to the risk in each

Contribution approach (CA)

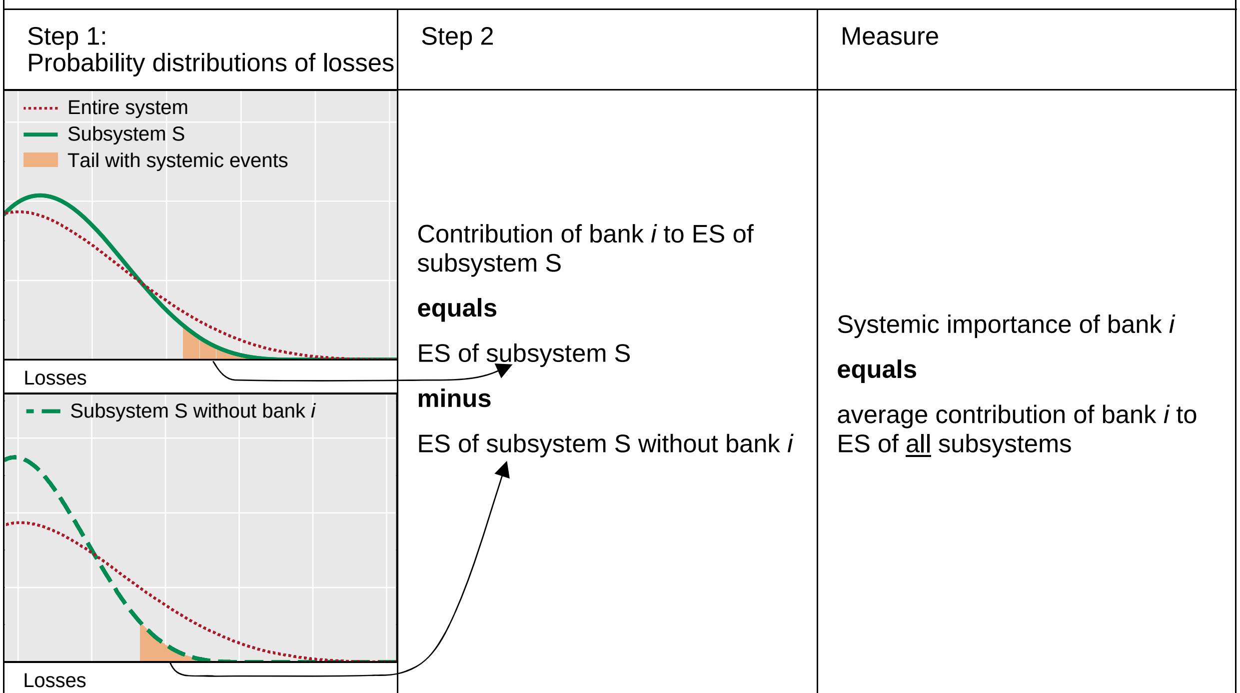

ES = expected shortfall. Graph 2

[4] The contribution approach was originally suggested by Tarashev et al (2010). It has been subsequently implemented by Gauthier et al (2010), Liu and Staum (2010) and Drehmann and Tarashev (2011) in a way that takes interbank links explicitly into account.

Bottom-up approach (BA)

Step 1	Step 2: Probability distribution of losses	Measure
Bank *i* defaults	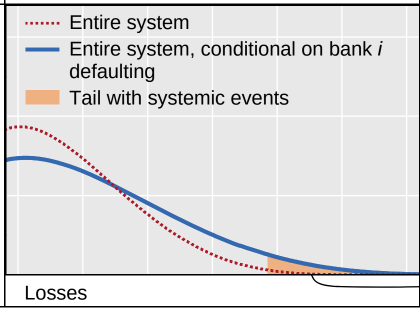	Systemic importance of bank *i* **equals** ES of entire system, conditional on bank *i* defaulting

ES = expected shortfall.

Graph 3

possible subsystem. The derivation of such a marginal contribution for a given subsystem *S* is illustrated in Graph 2.[5]

Bottom-up measure

The bottom-up approach (BA) reverses the logic of the PA. Namely, it measures the systemic importance of a bank by the ES of the whole system, conditional on this particular bank being in default. This is shown in Graph 3.[6]

Bottom-up measures have been implemented frequently in the literature. For example, conditional on the default of a bank, Elsinger et al (2006) measure the ES of all other banks, whereas Segoviano and Goodhart (2009) derive the probability that at least one more bank defaults. Similar measures have also been popular in network analysis (for an overview, see Allen and Babus (2009) or Upper (2011)). More recently, Adrian and Brunnermeier (2010) suggest using CoVaR, ie the system-wide value-at-risk (VaR), conditional on an institution being in distress.

Empirical setup

The empirical analysis …

There are two key building blocks of our empirical analysis. The first is a probability distribution of losses to banks' non-bank creditors, which is the basis of our measure of system-wide risk and each of the three alternative measures of individual banks' systemic importance. The second building block is a set of simple indicator variables that could proxy for the more sophisticated measures of systemic importance.

[5] For a more technical discussion, see Tarashev et al (2010).

[6] Our objective is to make the *conditional* ES, from the bottom-up measure, comparable to the *unconditional* ES at the 99% confidence level, which underpins the top-down measures. To meet this objective, we seek to align the systemic events over which the conditional and unconditional ES take expectations. It turns out that this is attained (on average across banks) for a conditional ES at the 75% confidence level.

Losses to non-bank creditors

Systemic risk in our setup, ie risk to non-bank creditors, stems exclusively from bank defaults. In turn, a default occurs if losses on a bank's assets wipe out its equity cushion. Such losses can arise from two sources. On the one hand, banks can experience losses on their non-bank exposures, which, if sufficiently large, trigger first-round defaults. On the other hand, credit losses on interbank exposures can cause additional bank failures, or second-round defaults.[7]

... employs data on 20 large banks ...

Several inputs, which we describe at some length in the box on page 31, play a key role in our derivation of first- and second-round defaults. One is the probability that each bank in our sample defaults. Our premise is that this probability is influenced by prudential rules that set capital requirements on the basis of bank-level information. The second input is data on the correlation of exogenous shocks. The higher this correlation, the more likely it is that, when one bank experiences a first-round default, other banks also default or have their balance sheets weakened. And a bank with a weaker balance sheet is more likely to experience a second-round default if it is exposed to a defaulted bank. We capture the size of interbank exposures through estimates of bilateral interbank positions.

... in order to compare the rigorous measures ...

We derive default probabilities, asset correlations and interbank exposures for a system of 20 large internationally active banks on the basis of data from 2006–09. Then, we simulate exogenous shocks to the claims of the banks in this system on non-banks. This ultimately delivers the joint probability distribution of losses to non-bank creditors. Based on this distribution, we derive the system's expected shortfall as our measure of systemic risk and implement the three alternative approaches to systemic importance as outlined above. For more detail on the implementation, see Drehmann and Tarashev (2011).

Simple indicators

It is unlikely that a prudential authority will *directly* employ any of the sophisticated measures delivered by the alternative approaches to systemic importance. Instead, the authority may derive these measures only to approximate them with simple and reliable bank-specific indicators. Basing actual policy on such indicators, authorities would trade precision of the assessment for transparency and ease of communication. Furthermore, authorities would also allow banks that do not have system-wide information to assess and manage their own systemic importance.

... to simple indicators of systemic importance

We examine the information in three indicators of systemic importance. All come directly from banks' financial statements. The first is *bank size*. Given our focus on non-bank creditors, we measure size as a bank's liabilities to non-banks. The other two indicators relate to linkages across banks. One of them is *interbank lending* (IL). This provides information on the degree to which a bank is exposed to risks stemming from the interbank market. The other one –

[7] "Second-round defaults" refers to failures induced through direct interlinkages but also failures resulting from longer domino-type default cascades.

Data and calibration[1]

We analyse a system of 20 large banks on the basis of two sets of data.[2] The first comprises estimated correlations of asset returns between 2006 and 2009. We use these estimates to generate correlated shocks to banks' claims on non-banks. The second dataset refers to banks' balance sheets at end-2009 (for our main analysis) and end-2006 (for a robustness check). We divide the assets side of each bank's balance sheet into interbank claims (precisely, loans and advances to banks) and claims on non-banks (total assets minus interbank claims). In turn, we divide the liabilities side into: interbank debt liabilities (deposits from banks), equity capital and debt liabilities to non-banks (total liabilities minus interbank debt liabilities minus equity capital).[3]

In order to simulate the probability distribution of losses in the system, we need information on each bank's probability of default (PD). We start with the premise that prudential authorities do not take a system-wide perspective. They set capital requirements based on bank-level information that does not reflect fully the complexity of counterparty exposures and system-level interbank linkages. In order to work in a straightforward setup, we then assume that each bank's probability of a first-round default is fixed at 0.1% but banks' different interbank exposures lead to different probabilities of second-round defaults and, thus, to different overall PDs.[4] We implement this assumption by adjusting the marginal probability distribution of the exogenous shocks to each bank's claims on non-banks.

For second-round defaults, we need information on the bilateral linkages across the 20 banks in our sample.[5] Since our data reveal only the total interbank positions on the balance sheet of each bank, we need to make certain assumptions. First, we assume that interbank linkages are fully captured by balance sheet data, thus excluding for instance securitised assets and derivative exposures. Second, we follow the literature in constructing a network of bilateral interbank linkages on the assumption that each bank in our sample spreads its entire interbank positions as evenly as possible across the other banks in the sample (Upper (2011)). Third, as in any empirical setting, our system is not truly closed in the sense that aggregate interbank assets are not exactly equal to aggregate interbank liabilities. Following the literature, we close the system by introducing a hypothetical "sink" bank.

The default of a bank, irrespective of whether it is first- or second-round, imposes losses on the bank's non-bank creditors. The magnitude of these losses depends on the level of the defaulted bank's assets and on bankruptcy costs. We assume that bankruptcy costs wipe out 20% of the bank's assets at default.

____

[1] For further detail, see Drehmann and Tarashev (2011). [2] These banks are: Bank of America, Barclays, BNP Paribas, Citigroup, Commerzbank, Crédit Agricole, Credit Suisse, Deutsche Bank, HSBC, ING, JPMorgan, Lloyds, Mizuho, Royal Bank of Scotland, Santander, Société Générale, Sumitomo Mitsui, UBS, UniCredit and Wells Fargo. [3] Our data sources are Moody's KMV and Bankscope. [4] As a robustness check, we assess the empirical performance of the indicators under the assumption that supervisors can control a bank's overall PD, stemming from first- and second-round defaults. This setup leaves our conclusions broadly unchanged. [5] In principle, the correlations of banks' asset returns reflect both common exposures to non-banks and interbank linkages. Background analysis reveals, however, that interbank linkages affect the tail of the distribution of asset returns and, thus, have a negligible impact on asset return correlations, which are related mainly to the centre of this distribution. We abstract from this impact in our calibration of the banking system.

interbank borrowing (IB) – captures the extent to which a bank imposes credit risk on other banks and, thus, can propagate shocks through the system.

Empirical results

In this section, we report and discuss the results of our empirical analysis. We start with a comparison of the three model-based measures of systemic importance. We then examine the performance of our simple indicators as proxies for the three measures.

Systemic importance of banks[1]

Relative to the average bank

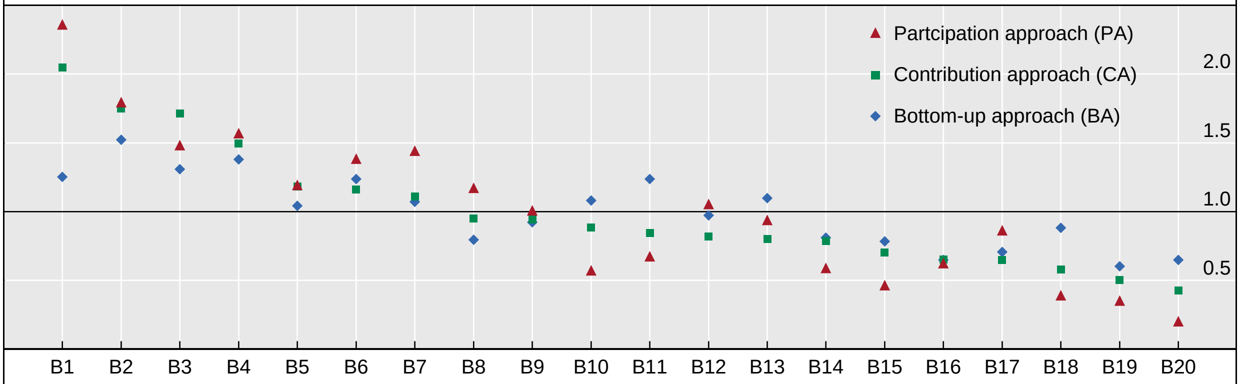

[1] Banks are ordered, along the horizontal axis, according to their level of systemic importance under CA. Graph 4

Differences across measures of systemic importance

Since they have different conceptual underpinnings, the alternative model-based measures of systemic importance can be expected to deliver different results. We examine differences across measures on the basis of the implied levels of systemic importance reported in Graph 4. In this graph, banks are ordered according to their systemic importance under CA (green squares), so that bank 1 is the most systemically important bank under this approach. The red triangles plot systemic importance under PA, and the blue diamonds under BA.

Indeed, there is a pronounced variation in the measured levels of systemic importance. On average, the absolute difference between the PA and CA measures is roughly 20% of CA measures. Furthermore, for nearly a third of the banks in our sample, these differences are between 30 and 50%. Similar discrepancies exist between BA and either PA or CA measures.

Despite differences across measures of systemic importance …

Such differences between alternative measures of systemic importance have important policy implications. For example, they indicate that capital requirements calibrated to the levels of systemic importance could depend materially on the approach chosen. This underscores that policymakers need to be careful in picking a measure that is aligned with their own perspective on systemic importance. But it also raises the question whether prudential policy should be based on just one measure or whether it should be guided by simple indicators. We turn to this question next.

Explanatory power of the simple indicators

Given the differences across the three approaches, how should we expect them to relate to our simple indicators? Under any approach, measured systemic importance should be expected to increase in size because the default of a bank that has borrowed more leads to larger losses. By contrast, the three measures of systemic importance are likely to differ in the way they relate to the interbank market indicators, IL and IB.

Drehmann and Tarashev (2011) discuss why interbank positions are captured differently by CA and PA. If a portion of systemic risk is associated with a particular interbank link, then CA splits this portion equally between the interbank lender and borrower. In this sense, CA provides a "fair" attribution of systemic risk. By contrast, PA focuses squarely on the losses that a bank could impose on *its own* non-bank creditors. All else equal, the risk of such losses is higher in the case of an interbank lender, who is vulnerable to shocks from the interbank market. Thus, PA attributes most of the risk of an interbank transaction to the lender.

BA treats interbank positions in the opposite way to PA. By considering systemic risk when a particular bank defaults, BA attributes a higher level of systemic importance to a bank whose default poses risk to other banks. This is

Systemic importance and simple indicators[1]

2009 balance sheets

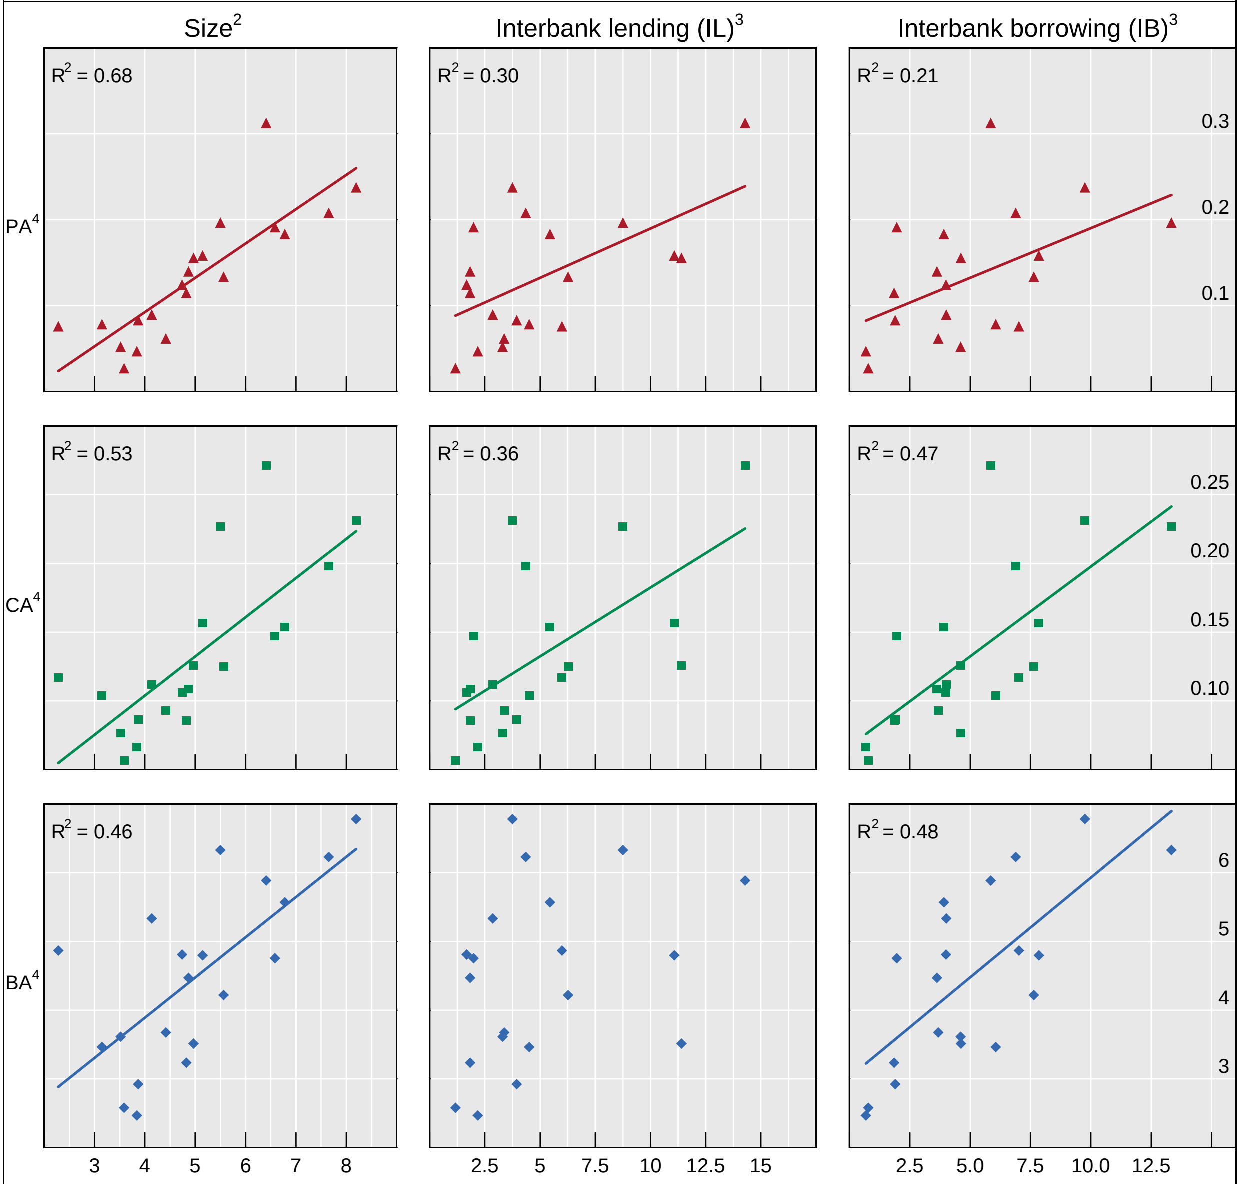

PA = participation approach; CA = contribution approach; BA = bottom-up approach.

[1] Lines plot the fitted values implied by bivariate linear regressions, given that the slope coefficient is different from zero with 95% confidence. [2] Non-bank debt liabilities of a bank relative to system aggregate, in per cent. Along the horizontal axis. [3] Relative to system aggregate, in per cent. Along the horizontal axis. [4] Systemic importance, per unit of system size, in per cent.

Graph 5

Regression results[1]												
	Participation approach (PA)				Contribution approach (CA)				Bottom-up approach (BA)			
2009 balance sheets												
Size	4.0***			3.5***	2.9***			2.1***	58.6***			43.4***
IL		10.4***		7.3***		9.0***		4.8***		106.6		–20.3
IB			10.6**	0.8			12.0***	5.8***			267.4***	214.4**
R^2	0.68	0.30	0.21	0.85	0.53	0.36	0.47	0.84	0.46	0.07	0.48	0.70
2006 balance sheets												
Size	5.1***			3.9***	4.0***			1.9***	67.8***			53.9***
IL		15.1***		9.5***		18.5***		12.4***		177.8**		97.3
IB			15.9***	1.8			18.8***	6.1*			198.7***	30.1
R^2	0.66	0.50	0.48	0.89	0.34	0.72	0.64	0.86	0.45	0.25	0.28	0.51

The dependent variable, systemic importance under alternative approaches, is in tens of millions of US dollars. Size = non-bank liabilities; IL (interbank lending) = total interbank assets; IB (interbank borrowing) = total interbank liabilities (all in billions of US dollars).

[1] Constants are included in all regressions but not reported for brevity. ***/**/* indicates that the coefficient is significant at the 1%/5%/10% level, respectively. Adjusted R^2s are reported. Table 1

the case for a large interbank borrower. As a result, BA assigns a higher level of systemic importance to this bank than to another similarly sized bank which is primarily an interbank lender.

In summary, we would expect the following relationship between the three simple indicators and the three measures of systemic importance. First, size should be important for all measures. Second, interbank lending (IL) should have strong explanatory power for measures obtained under CA and PA, whereas interbank borrowing (IB) should help to predict CA and BA.

... bank size helps explain each of them

Graph 5 visualises the relationship between the simple indicators and the model-based measures of systemic importance. A row in this graph corresponds to a particular measure (PA, CA or BA) and a column to a particular indicator (size, IL or IB). The bivariate regression lines show the estimated relationship between a simple indicator and a more rigorous measure, if it is statistically significant at the 95% confidence level (all regression results are reported in Table 1).

As expected, size (first row) is a robust indicator across all measures. And it consistently exhibits the highest explanatory power (as captured by the goodness of fit statistic R^2). The effects of interbank lending (second row) are also in line with the earlier discussion: a larger IL is associated with a higher level of systemic importance under PA and CA but not under BA. Likewise, being a large interbank borrower typically leads to high levels of measured systemic importance (third row).

Indicators of interbank activity ...

Graph 5 also highlights that different indicators carry complementary information about systemic importance. For instance, the bank with the highest level of systemic importance under PA and CA is not the largest bank in the sample. But it is the most active lender in the interbank market. Similarly, the seventh largest bank attains the third highest level of systemic importance under CA as well as BA because it is the largest interbank market borrower.

... provide complementary information

Simple multivariate regressions indicate the degree to which different indicators carry complementary information (Table 1, upper panels).[8] Size remains statistically significant once all indicator variables are included. By contrast, the performance of interbank borrowing and lending (IB and IL) does depend on the approach underlying the model-based measure of systemic importance. Fully in line with the above discussion, IB is not an important driver of systemic importance under PA, which attributes the risk associated with an interbank link mainly to the lending counterparty. Likewise, since BA assigns this risk mainly to the interbank borrower, it tends to render IL uninformative. Only under CA are borrowers and lenders treated equally. Thus, both IL and IB help to explain CA measures of systemic importance.

In order to verify the robustness of the above results, we rerun the linear regressions after recalculating the alternative measures and indicators on the basis of 2006, instead of 2009, balance sheet data. The new results (reported in the lower panels in Table 1) confirm our previous conclusions about the explanatory power of size as well as the lending and borrowing indicators.[9]

Economic significance of the simple indicators

Indicators' economic significance ...

What do these results mean for the economic significance of each indicator? This is not fully apparent from the regression results in Table 1, as the top-down and bottom-up approaches measure systemic importance from different perspectives, thus impairing the comparability of regression coefficients. For each approach, we therefore calculate the predicted level of

Economic impact[1]

	PA	CA	BA	Average
Size	12.24	6.14	4.88	7.75
IL	3.33	2.85	0.00	2.06
IB	0.00	2.27	1.27	1.18

PA = participation approach; CA = contribution approach; BA = bottom-up approach; IL = interbank lending; IB = interbank borrowing. For definitions, see Table 1.

[1] Economic impact is measured by the ratio of predicted systemic importance if an indicator is increased by 10% relative to the predicted level of systemic importance for a bank with average size and the average level of interbank lending and interbank borrowing. Predicted levels of systemic importance are based on the regression results shown in Table 1, setting insignificant coefficients to zero and averaging across 2006 and 2009 results. In per cent. Table 2

[8] We consider linear regressions in order to study the explanatory power of the indicators under a simple specification. In general, the true relationship between an indicator and a measure of systemic importance would be non-linear. Tarashev et al (2010) derive this formally in the case of bank size.

[9] The weakening of IB's explanatory power for CA and BA is the result of multicollinearity. IB exhibits significant explanatory power on a standalone basis. However, it loses its explanatory power in a regression with all three indicators because it is highly correlated with IL (a correlation coefficient of 71%).

systemic importance for a bank of average size and average levels of interbank lending and interbank borrowing.[10] Then, we increase each indicator by 10% and recalculate the predicted level of systemic importance under PA, CA and BA (Table 2).

Taking the results at face value, we conclude that the economic significance of size is much larger than that of the other two indicators. Concretely, increasing size by 10% has a two to four times greater impact on systemic importance than increasing IL or IB by 10%. In turn, it seems that IL is economically more significant than IB.

… differs across measures of systemic importance

Table 2 also shows that the economic significance of each indicator depends materially on the measure of systemic importance. As in the case of regression results, this is in line with the different economic logic underlying the three measures of systemic importance. The policy implication of the finding is that, even when simple indicators are used, regulators should have a clear understanding of their preferred perspective on systemic importance.

Conclusion

In this article, we investigate whether simple indicators can approximate more complex measures of a bank's systemic importance. And since systemic importance itself is a multifaceted concept, we measure it from three different perspectives, based on top-down or bottom-up approaches.

We find that bank size is a reliable proxy of systemic importance, regardless of the perspective chosen. Interbank lending or borrowing provides additional useful information for some measures but not for others. This result is not surprising as it is fully in line with the economic logic underlying each measurement approach.

Taken together, our results highlight that simple indicators do help to assess the degree of banks' systemic importance. Given the complexities of implementing and communicating more rigorous model-based measures of systemic importance, these results suggest that an indicator approach may be the most suitable route for practical purposes. It would also allow banks with limited system-level information to measure and manage their own systemic importance.

[10] For this calculation, we set all insignificant coefficients to zero and average across the regression results of 2006 and 2009.

References

Acharya V, L Pedersen, T Philippon and M Richardson (2009): "Measuring systemic risk", New York University, mimeo.

Adrian, T and M Brunnermeier (2010): "CoVaR", *Working Paper*, Princeton University.

Allen, F and A Babus (2009): "Networks in finance", in P Kleindorfer and J Wind (eds), *The network challenge*, Wharton School Publishing, pp 367–80.

Brownlees, C and R Engle (2010): "Volatility, correlation and tails for systemic risk measurement", New York University, mimeo.

Drehmann, M and N Tarashev (2011): "Measuring the systemic importance of interconnected banks", *BIS Working Papers*, no 342.

Elsinger, H, A Lehar and M Summer (2006): "Using market information for banking system risk assessment", *International Journal of Central Banking*, 2, pp 137–65.

Financial Stability Board, International Monetary Fund, Bank for International Settlements (2009): "Guidance to assess the systemic importance of financial institutions, markets and instruments: initial considerations", *FSB Background Paper*.

Gauthier C, A Lehar and M Souissi (2010): "Macroprudential regulation and systemic capital requirements", *Working Papers*, 2010-4, Bank of Canada.

Huang X, H Zhou and H Zhu (2010): "Assessing the systemic risk of a heterogeneous portfolio of banks during the recent crisis", *BIS Working Papers*, no 296.

Liu, M and J Staum (2010): "Systemic risk components in a network model of contagion", mimeo, Northwestern University.

Segoviano, M and C Goodhart (2009): "Banking stability measures", *IMF Working Paper*, WP/09/4 0020.

Shapley, L (1953): "A value for n-person games", in H Kuhn and A Tucker (eds), *Contributions to the theory of games*, volume II, *Annals of Mathematical Studies*, v 28, pp 307–17, Princeton University Press.

Tarashev, N, C Borio and K Tsatsaronis (2010): "Attributing systemic risk to individual institutions", *BIS Working papers*, no 308.

Upper, C (2011): "Simulation methods to assess the danger of contagion in interbank markets", *Journal of Financial Stability*, forthcoming.

Webber, L and M Willison (2011): "Systemic capital requirements", *Bank of England Working Papers*, forthcoming.