Nikola Tarashev
nikola.tarashev@bis.org

Claudio Borio
claudio.borio@bis.org

Kostas Tsatsaronis
ktsatsaronis@bis.org

The systemic importance of financial institutions[1]

Prudential tools that target financial stability need to be calibrated at the level of the financial system but implemented at the level of each regulated institution. They require a methodology for the allocation of system-wide risk to the individual institution in line with its systemic importance. This article proposes a general and flexible allocation methodology and uses it to identify and quantify the drivers of systemic importance. It then illustrates how the methodology could be employed in practice, based on a sample of large internationally active institutions.

JEL Classification: C15, C71, G20, G28.

On 16 September 2008 the US authorities announced that they would take the unprecedented step of offering emergency financial support to AIG, a large insurance conglomerate. The decision was rooted in concerns about the repercussions of the failure of this institution on the economy at large, ie about its systemic importance.[2] Similar far-reaching and urgent decisions were taken by authorities in other jurisdictions. By contrast, in 1995, the Bank of England had allowed merchant bank Barings to fail because it considered this would have no material impact on other banks (which was subsequently confirmed).

More generally, the events of the past two years serve as a stark reminder that systemic financial disruptions can have large macroeconomic effects. As a result, the objective of strengthening the macroprudential orientation of financial stability frameworks has risen to the top of the international agenda.[3] The main distinction between the macro- and microprudential perspectives is that the former focuses on the financial system as a whole, whereas the latter focuses on individual institutions.[4]

[1] The authors thank Marek Hlavacek for excellent research assistance, and Stephen Cecchetti, Robert McCauley and Christian Upper for helpful comments. The views expressed in this article are those of the authors and not necessarily those of the BIS.

[2] The press release from the Federal Reserve explained: “The Board determined that, in current circumstances, the disorderly failure of AIG could add to already significant levels of financial market fragility and lead to substantially higher borrowing costs, reduced household wealth, and materially weaker economic performance.”

[3] See G20 (2009) and de Larosiere (2009) for reports on this international consensus.

[4] See Crockett (2000), Knight (2006) and Borio (2009) for an elaboration of the macroprudential approach and progress in its implementation.

By necessity, however, the tools of financial sector supervision and key policy interventions are applied to individual institutions, even when decisions are motivated by systemic considerations. Thus, policymakers need analytical tools to help them assess the systemic importance of individual institutions. In times of crisis, these tools can help to gauge the likely impact of distress at a given financial firm on the stability of the overall financial system. In periods of calm, they can help to calibrate prudential instruments, such as capital requirements and insurance premiums, according to the relative contribution of different institutions to systemic risk.

This article presents a methodology that takes as inputs measures of system-wide risk and allocates them to individual institutions. The methodology is derived directly from a game-theoretic concept, the Shapley value, which describes a way of allocating the collective benefit created by a group to the individual contributors. The Shapley value approach satisfies a number of intuitive criteria and is quite general, being applicable to a wide spectrum of measures of system-wide risk.

The methodology makes it straightforward to quantify the impact of the various drivers of an institution's systemic importance. These include their riskiness on a standalone basis, their exposure to common risk factors and the degree of size concentration in the system. A key result is that the contribution of an institution to system-wide risk generally increases more than proportionately with its size.

We apply the methodology to real-world data on a sample of 20 large internationally active financial institutions. The results highlight the interaction among the various drivers of systemic importance. In our sample, none of them, taken in isolation, is a fully satisfactory proxy for systemic importance.

The article is organised in four sections. The first section describes the allocation procedure and its properties. The second section applies the procedure to a specific measure of systemic risk in hypothetical and highly stylised financial systems in order to analyse the impact of different drivers of systemic importance. The third section discusses how the methodology could be used in practice as a tool to mitigate systemic risk and applies it to real-world data. The last section concludes.

The allocation procedure: measuring systemic importance

A general methodology to attribute systemic risk ...

The problem of allocating system-wide risk to individual institutions is analogous to that of a risk controller in an investment firm seeking to attribute the use of the firm's risk capital to individual desk traders. The fact that the sum of the risks incurred by each desk in isolation does *not equal* the total risk for the firm complicates the controller's problem. Simple summation ignores that the interactions among individual positions could reduce or compound overall risk. They would reduce it when positions across desks partially cancel each other out; they would compound it when losses in one side of the business are incurred simultaneously with, or trigger, losses in another.

Game theorists have tackled similar problems in the context of cooperative games. These are general settings where a group of players engage in a collective effort in order to generate a shared benefit[5] (called "value") for the group. The theoretical problem of allocating this value among individual players in a way that satisfies certain fundamental criteria is conceptually identical to that of risk attribution described above.

... draws on game theory ...

Lloyd Shapley proposed a methodology that distributes the overall value among players on the basis of their individual contributions (Shapley (1953)). The idea behind the allocation methodology is quite simple. Adding up what individual players can achieve by themselves (the equivalent of summing up the standalone risk of each trading desk in the investment firm) is unlikely to reflect their contributions to the productivity of others. Similarly, calculating the marginal contribution of a single player as the difference between what the entire group can achieve with and without the specific individual gives only a partial picture of the individual's contribution to the work of others. The reason is that this method also ignores the complexities of bilateral relationships. By contrast, the Shapley methodology accounts fully for the degree to which such relationships affect the overall outcome. It accomplishes this by ascribing to individual players the average marginal contribution each makes to each

Box 1: Shapley value allocation methodology: a specific example

This box illustrates the Shapley value allocation methodology by reference to a specific numerical example where three parties (A, B and C) can cooperate to generate a measurable outcome. If nobody participates nothing is produced, and each participant alone can produce 4 units. The output of each possible grouping of the three participants is detailed in the left-hand column of the table below.

Subgroup	Subgroup output	Marginal contribution of A	Marginal contribution of B	Marginal contribution of C
A	4	4	.	.
B	4	.	4	.
C	4	.	.	4
A, B	9	5	5	.
A, C	10	6	.	6
B, C	11	.	7	7
A, B, C	15	4	5	6
Shapley value	.	4.5	5	5.5

Table A

The marginal contribution of a player to a subgroup is calculated as the output of the subgroup minus the output of the same subgroup excluding the individual participant. For instance, the marginal contribution of A to the output of the overall group (A, B, C) is equal to the difference between 15, which is the overall group's output, and 11, which is the output of B and C together.

The Shapley value of each player is the average of its marginal contributions across all differently sized subgroups. For example, the value of B is equal to 5 (see bottom row). It is calculated as the average of 4, which is its individual output, 6, which is the mean contribution it makes to subgroups of size two, and 5, which is its marginal contribution to the overall group. The calculation can also be motivated as the expected marginal contribution of an individual participant in groups that are formed randomly by sequentially selecting players (see Mas-Colell et al (1995)).

[5] This is a very general concept that could be thought of as wealth, or collective output.

possible subgroup in which they participate (see Box 1 on the previous page for a detailed exposition of the methodology and a numerical example).

... and has many appealing features:

In addition to its simplicity, the Shapley value has a number of intuitively appealing features.[6] It ensures that the gains from cooperation between *any two* players are divided equally between them; in other words, it is "fair" in the sense that it does not lead to biased outcomes that favour or penalise particular players. It distributes *exactly* the total benefit to all players, without resulting in any surplus or deficit. It is symmetric, in the sense that two players with the same characteristics receive the same share of the overall value. And it assigns no payoff to a player who makes no contribution to any subgroup.

it measures individual contributions to risk;

An application of the Shapley value methodology to the measurement of institutions' systemic importance simply transposes the problem of distributing a collective value among individual players to that of attributing overall risk to individual institutions. It requires as an input a quantitative measure of risk for all groupings of institutions. These range from the largest group comprising all institutions to the smallest, which consist of single institutions. The methodology then attributes the overall (system-wide) risk to each institution on the basis of its average contribution to the risk of all the groups in which it participates. *The degree of systemic importance of institutions is therefore captured by the share of systemic risk that is attributed to each of them.* Institutions with higher systemic importance will have a higher Shapley value than others.

is general and flexible;

A major strength of the Shapley value methodology is its generality. It accommodates any systemic risk measure that treats the system as a portfolio of institutions and identifies risk with the uncertainty about the returns (losses) on this portfolio. In addition, existing allocation procedures are specific applications of the Shapley value methodology. This is the case, for instance, of the procedure recently proposed by Acharya and Richardson (2009) for the calibration of institution-specific premiums for insurance against systemic distress. Tarashev et al (2009) discuss these points at some length.

and is robust to uncertainty

Another strength of the Shapley value methodology is that it allows measures of systemic importance to account for model and parameter uncertainty. Such uncertainty may make it natural to measure systemic risk under alternative models and parameter estimates. This would lead to alternative measures of systemic importance for each institution. Being linear, the Shapley value implies that the weighted average of alternative measures (a linear combination) can be used as a single robust measure of systemic importance.

6 For a fuller discussion of the technical properties of the Shapley value, see Mas-Colell et al (1995). Tarashev et al (2009) provide a more detailed description of how to implement a Shapley value decomposition in the context of the attribution of system-wide risk to individual institutions.

Drivers of systemic importance: stylised examples

Drivers of systemic risk

In this section we study three drivers of systemic risk and, hence, of the systemic importance of individual institutions. One is the riskiness of individual firms, as captured by their probabilities of default (PDs).[7] Another is the degree of size concentration, or "lumpiness", of the system, which increases as the number of institutions decreases or as their relative sizes become more disparate. The final driver is the institutions' exposure to common (or systematic) risk factors, which arises either because financial institutions are similar to each other (eg lend to the same sectors) or because they are interconnected. Importantly, while the probability of default (or insolvency) can be constructed on the basis of institution-specific characteristics alone, the other two drivers relate to characteristics of the system as a whole.

As a concrete measure of systemic risk, we use expected shortfall, which equals the expected (average) size of losses in a systemic event (see the appendix on page 86 for detail). In general, a systemic event is defined as one that generates losses deemed large enough to cause disruptions to the functioning of the system. In this article, a systemic event is defined as the occurrence of extreme aggregate losses that materialise with a given small probability, ie losses that exceed a certain threshold.[8]

The impact of the three drivers on systemic risk is quite intuitive. Keeping everything else constant, an increase in institutions' PDs leads to a higher level of systemic risk. Even if the PDs remain unchanged, greater lumpiness of the system reduces diversification benefits, raising the likelihood of extreme losses and, with it, expected shortfall. Similarly, greater exposure to common risk factors increases the likelihood of joint failures and hence also the likelihood of extreme losses in the system.

To explore the impact of the same three drivers on the systemic importance of individual institutions, we resort to numerical exercises. For these exercises, we allocate system-wide expected shortfall to individual institutions ("banks") on the basis of the Shapley value methodology. The results, based on highly stylised hypothetical systems, yield four key messages.

Systemic importance increases with …

First, a rise in an institution's exposure to a common risk factor increases its systemic importance. This is illustrated in Table 1, which compares a number of banking systems, each comprising 20 banks. In every system there are two homogeneous groups, A and B, which differ only with respect to banks' exposures to the common factor. Keeping the strength of exposures to the common factor in group B constant but increasing it for group A (across columns to the right, in each panel) results in an increase in these banks' share in systemic risk. In the specific example of a strongly capitalised system, the

---

[7] Strictly speaking, an institution's standalone risk depends both on its PD and on its loss-given-default (LGD). This article abstracts from LGD by assuming that it is constant and equal for all financial institutions. Relaxing this assumption in order to account for certain empirical properties of LGD would not alter any of the qualitative conclusions derived below.

[8] A similar setting has been used in the context of financial stability by Kuritzkes et al (2005), who measure the expected loss to the deposit insurance fund using similar concepts.

Common exposures, systemic risk and systemic importance										
	Strongly capitalised system (all PDs = 0.1%)					Weakly capitalised system (all PDs = 0.3%)				
	Exposure to the systematic risk factor (banks in group A)					Exposure to the systematic risk factor (banks in group A)				
	ρ = 0.30	ρ = 0.40	ρ = 0.50	ρ = 0.60	ρ = 0.70	ρ = 0.30	ρ = 0.40	ρ = 0.50	ρ = 0.60	ρ = 0.70
Group A (share)	44.0%	46.2%	50.0%	54.4%	60.4%	41.7%	45.4%	50.0%	56.2%	63.2%
Group B (share)	56.0%	53.8%	50.0%	45.6%	39.6%	58.3%	54.6%	50.0%	43.8%	36.8%
Total ES	4.0	4.4	5.0	5.8	6.8	6.6	7.2	8.2	9.8	11.5

Total expected shortfall (ES) equals the expected loss in the 0.2% right-hand tail of the distribution of portfolio losses; per unit of overall system size, in percentage points. The first two rows report the share of the two groups (each comprising 10 banks) in total ES. The exposure of each of the 10 banks in group A to the systematic risk factor is as given in the row headings. The exposure of each of the 10 banks in group B to the systematic risk factor corresponds to ρ = 0.50. See the technical appendix for a definition of ρ. The probability of defaut (PD) of each bank is as specified in the panel heading. Loss-given-default is set to 55%. All banks are of equal size, each one accounting for 5% of the overall size of the system. Table 1

combined contribution of group A banks rises from 44% to roughly 60%. The result is similar for a weakly capitalised system.

... the strength of common risk factors ...

The reason for this result is that higher exposures to the common factor result in a higher probability of joint failures in the system. In turn, a higher probability of joint failures translates into higher average losses in the systemic event, which leads to a higher level of systemic risk, as measured by expected shortfall. Quite intuitively, the rise in the level of systemic risk is attributed mainly to the banks that contribute most to this rise, ie those that experience an increase in their exposure to the common factor (group A banks in Table 1).

... individual riskiness ...

Second, the interaction between different drivers may reinforce their impact on systemic importance. A concrete example is provided by Graph 1 (left-hand panel) on the basis of a system in which banks differ only in terms of size. As the strength of exposures to the common factor increases uniformly across all banks in this system, the portion of the expected shortfall

Systemic risk: interaction of different drivers[1]

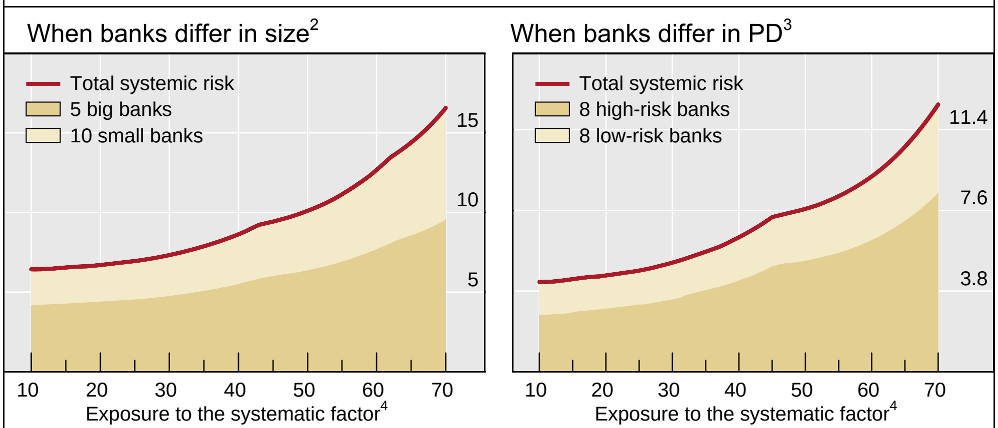

[1] All numbers are in percentage points. Total systemic risk equals the expected loss in the 0.2% right-hand tail of the distribution of portfolio losses; per unit of overall system size. The contributions of the two groups of banks to the total are plotted as shaded areas. Each group accounts for half of the overall system size. Loss-given-default is assumed to be 55%. [2] Each bank's probability of default (PD) equals 0.3%. [3] The PD of a high-risk bank is 0.3% and that of a low-risk bank is 0.1%. [4] See the technical appendix for a definition. Graph 1

attributable to larger banks increases by a greater amount than that attributable to smaller banks. In other words, bank size reinforces the impact of common factor exposures on systemic importance. The right-hand panel of Graph 1 illustrates a similar point in the context of a system comprising banks that differ only with respect to their individual PDs. If all of these banks experience the same rise in their exposures to the common factor, the increase in the contributions to systemic risk is greater for riskier banks. Here, individual riskiness reinforces the impact of common factor exposures on systemic importance.

... and institutions' relative size

Third, changing the lumpiness of a system affects the systemic contributions of banks of different sizes differently. This is reported in Table 2, which considers hypothetical banking systems where all banks feature the same PDs and exposures to the common factor but differ in size. There are three big banks of equal size, together accounting for 40% of the overall system, and a group of small banks, making up the rest. As the number (but not the share) of small banks increases (across columns to the right, in each panel), diversification benefits reduce overall systemic risk.[9] This reduction is associated with a decline in the systemic importance of small banks and a rise in that of large banks (the first two rows in each panel). Moreover, the rise in big banks' systemic importance reflects not only a rise in the share but also in the amount of systemic risk that these banks account for. Considering the example of a strongly capitalised system (left-hand panel), a rise in the number of small banks from five to 25 results in a drop of systemic risk from 9.8 to 9.3 cents on the dollar. At the same time, the amount of this risk that big banks account for rises from 4.3 (or 42.8% of 9.8) to 6.3 (or 68.1% of 9.3) cents on the dollar.[10]

System lumpiness, systemic risk and systemic importance

	Strongly capitalised system (all PDs = 0.1%)					Weakly capitalised system (all PDs = 0.3%)				
	Number of small banks					Number of small banks				
	n_s = 5	n_s = 10	n_s = 15	n_s = 20	n_s = 25	n_s = 5	n_s = 10	n_s = 15	n_s = 20	n_s = 25
Three big banks (share)	42.8%	56.8%	62.6%	66.0%	68.1%	41.6%	52.3%	56.5%	59.3%	60.7%
n_s small banks (share)	57.2%	43.2%	37.4%	34.0%	31.9%	58.4%	47.7%	43.5%	40.7%	39.3%
Total ES	9.8	9.4	9.3	9.25	9.23	16.7	15.0	14.7	14.4	14.3

Total expected shortfall (ES) equals the expected loss in the 0.2% right-hand tail of the distribution of portfolio losses; per unit of overall system size, in percentage points. The first two rows report the share of the two groups of banks in total ES. The group of big banks accounts for 40% of the overall size of the system and the group of small banks accounts for 60%. The probability of default (PD) of each bank is as specified in the panel heading. Loss-given-default is set to 55%. All banks are assumed to have the same sensitivity to common risk factors, implying a common asset return correlation of 42%. Table 2

[9] The decline in systemic risk is rather subdued because the assumed high exposure of banks to the common risk factor restricts the diversification benefits obtained from increasing their number. This general result is studied in detail in Tarashev (2009).

[10] The effect is even stronger in the case of a weakly capitalised system (right-hand panel).

Size and systemic importance[1]

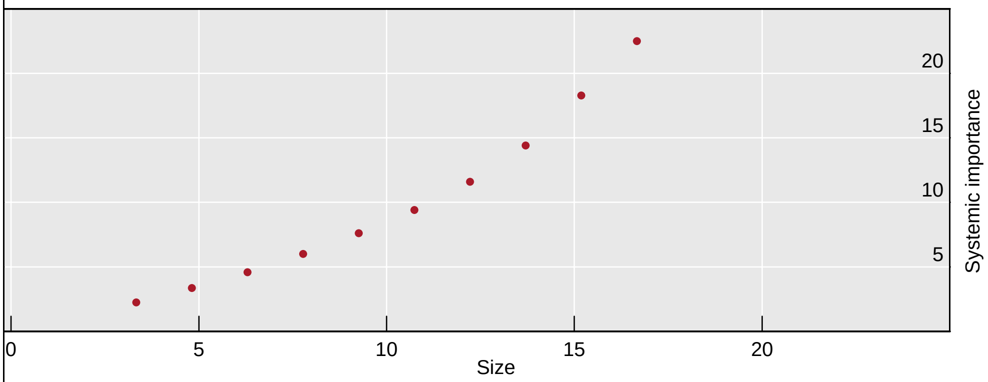

[1] All numbers are in percentage points. The system comprises 10 institutions, each represented by a dot. Systemic importance is measured as the share of each institution in the expected shortfall of the system, defined as the expected loss in the 0.2% right-hand tail of the distribution of system-wide losses. Size is measured as a share in the aggregate size of all institutions in the system. Each bank's loss-given-default and probability of default equal 55% and 0.1% respectively. The loadings on the common factor (see the technical appendix) are constant across banks and equal 0.6. Graph 2

Finally, and quite generally, systemic importance increases more than proportionately with (relative) size. This relationship is a consequence of the fact that larger institutions play a disproportionate role in systemic events. The first column of Table 2, for example, relates to a system in which a big bank is roughly 10% larger than a small one but is assigned a 25% greater share in systemic risk.[11] This effect increases as banks' sizes become more disparate. For example, the fifth column of the table, which relates to a system where the sizes of big and small banks are roughly 5:1, reports that the respective shares in systemic risk are roughly 18:1.

Graph 2 presents further evidence of this non-linear relationship between size and systemic importance. It plots the contributions to system-wide risk of institutions that are all identical except for their size. In the particular example, the largest institution is about 5 times as large as the smallest one, but its relative systemic importance is nearly 10 times as high.

Implications for the calibration of prudential instruments

Even though the above examples have been cast in stylised settings, they illustrate robust results and point to concrete policy lessons. In particular, all else equal, they suggest that any "systemic capital charge" applied to individual institutions should increase more than proportionately with relative size. In other words, there is a clear rationale for having tighter prudential standards for larger institutions. In addition, the charge should increase with the degree to which an institution is exposed to sources of systematic risk. This means that higher capital charges would be applied to institutions that are more similar to the typical (or "average") institution: if they fail, they are more likely to fail in a systemic event.

The above examples also touch, albeit indirectly, on the notion of diversification from a systemic viewpoint. There is a potential trade-off between diversification in the portfolio of an *individual* institution and diversification *for*

[11] More precisely, the ratio of small and big bank sizes equals (0.4/3)/(0.6/5) = 1.11. The corresponding ratio of systemic importance measures is (42.8%/3)/(57.2%/5) = 1.25.

the system as a whole. This is because, by diversifying their own investment portfolios, institutions affect systemic risk in two ways. First, greater diversification of each portfolio is likely to reduce the riskiness of individual institutions. Second, it is also likely to result in more similar portfolios and, thus, in institutions being more exposed to common risk factors. The net outcome depends on how the first effect, which lowers systemic risk, compares to the second, which raises it.

Implementing the tool: beyond stylised examples

The previous analysis provides a structured framework for examining what factors are relevant in assessing the systemic importance of institutions. But what steps are needed to apply the Shapley value methodology in practice? What choices do policymakers have to make?

Operationalising the methodology

In making this general approach operational, a number of issues need to be addressed. Beyond choosing a specific measure of systemic risk, these include: the definition of the relevant "system"; the definition of the "size" of institutions; the choice of inputs; the uncertainty about the correct specification of the risk model and the true parameter values; and computational burden. Except for the last, all of these issues are related to the measure of systemic risk, rather than to the Shapley value methodology as such. Box 2 provides a discussion of the trade-offs and pitfalls involved and outlines the considerations that might guide policymakers' choices.

An application to real-world data …

Once these choices are made, the application is straightforward. To illustrate how the methodology can be applied to real-world data, consider the following example. The chosen measure of system-wide risk is expected shortfall, as in the stylised examples of the previous section. We define the relevant "system" as comprising 20 large internationally active financial institutions and assume that a loss is incurred when one or more of them fail. We measure an institution's size as the book value of its liabilities, divided by

A system of large internationally active institutions[1]

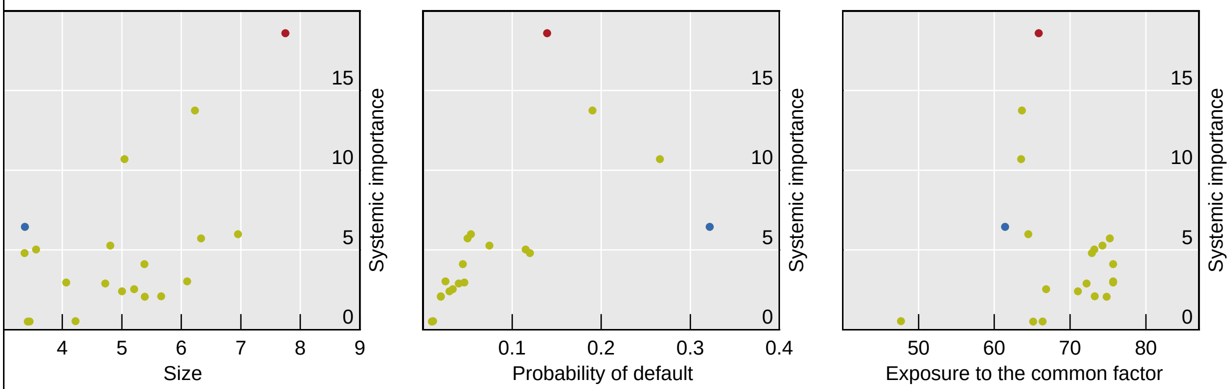

[1] All numbers are in percentage points. Systemic importance is measured as the share of each institution in the expected shortfall of the system, which is defined as the expected loss in the 0.2% right-hand tail of the distribution of portfolio losses. The size of an institution equals the book value of its liabilities, expressed as a share in the sum of the liabilities of all institutions in the system. The probability of default is the one-year EDF provided by Moody's KMV for end-2007. Exposures to the common factor are derived on the basis of Moody's KMV GCorr estimates of institutions' asset-return correlations for end-2007.

Sources: Moody's KMV. Graph 3

Box 2: Applying the method in a policy context: choices and trade-offs

This box addresses the policy choices and practical issues that have to be confronted when implementing the methodology as an element in a macroprudential approach to regulation and supervision.

The definition of the appropriate "system", as a precondition for calibration, is not straightforward. This is less of an issue in current regulatory arrangements which focus on individual institutions but becomes critical when the prudential framework focuses on systemic risk. At least two aspects need to be addressed. The first relates to the institutional coverage of regulation – its so-called "perimeter". A systemic approach would need to take account of the risks generated by all financial institutions that are capable, on their own and as a group, of causing material system-wide damage. This is so regardless of their legal form. The second aspect relates to the geographical coverage of regulation. Should the approach be applied at a domestic level or at a more global level, say to internationally active institutions? And if the answer is to both, how would the adjustments be reconciled? Clearly, a large dose of pragmatism is necessary. And the precise answers will also depend on the extent of cooperation across regulatory jurisdictions.

The definition of the size of the institutions also merits attention, and partly overlaps with that of the system. One question is whether to include only domestic exposures or both domestic and international ones. Another question is whether the appropriate measure refers to the assets (presumably including off-balance sheet items) or to the liabilities (excluding equity) of the institutions. Total assets better reflect the potential overall losses incurred by all the claimants on the institution; liabilities are a better measure of the direct losses linked to its failure.

Having defined the system and the size of the institutions, the next practical question is how to estimate the additional parameters, notably the probabilities of default and the factor loadings on the systematic risk factors. The sources of information range from market inputs, at one end, to supervisory inputs, at the other. Combinations of the two are also possible.

Market inputs have a number of attractive features but also limitations. On the plus side: they summarise the considered opinion of market participants based on the information at their disposal; they should reflect market participants' views of all potential sources of risk, regardless of their origin (eg poor asset quality, bank runs, counterparty linkages); and they are easily available on a timely basis. On the minus side: they may not be available for all institutions (eg equity prices for savings banks); they require the use of "models" to either filter out extraneous information (eg risk premia, expectations of bailouts) or complete the information they contain (eg to derive probabilities of default from equity prices), giving rise to "model" uncertainty; and they may contain systematic biases: for example, it is well known that market prices tend to be especially buoyant as financial vulnerabilities build up during booms (Borio and Drehmann (2009)).

Supervisory estimates have their own strengths and weaknesses. On the plus side, they can be based on more granular and private information, to which market participants do not have access; on the minus side, they may simply not be available, or may be hard to construct for certain inputs. For example, supervisors have a long tradition in producing measures of the soundness of individual financial institutions, such as rating systems. However, they have as yet not developed tools to derive measures of exposures to systematic risk factors and correlations across institutions based on balance sheet data. The available techniques are in their early stages of development.

All this suggests that, in practice, it might be helpful to rely on a combination of sources and to minimise their individual limitations. For example, currently market prices appear to be especially suited for the estimation of exposures to common factors. And long-term averages of such prices would help to address the biases in the time dimension. This would be especially appropriate if the tool is used to calculate *relative* contributions of institutions to systemic risk and to avoid procyclicality (Borio (2009)).

These difficulties highlight the need to deal with the margin of error that will inevitably surround the estimates of systemic risk and hence, by implication, of institutions' contributions to it (Tarashev (2009)). Fortunately, as noted above, the linearity property of the allocation procedure makes it possible to address this issue in a formal, simple and transparent way. This property allows one to combine alternative estimates, weighting them by the degree of confidence that one attaches to them (Tarashev et al (2009)). In addition, it may be advisable for policymakers not to rely too heavily on the resulting point estimates. One possibility would be to allocate institutions into a few buckets, each of them comprising an interval of point estimates – akin to a rating system. This grouping has the added advantage of reducing the computational burden of assessing risk at the level of subgroups of institutions.

the sum of the liabilities of all institutions in the system. In addition, we measure an institution's standalone riskiness as the Moody's KMV estimate of its one-year probability of default and assume that loss-given-default is constant at 55%. We also impose a single-common-factor structure on the Moody's KMV estimate of the 20 institutions' asset-return correlations in order to derive the strength of exposures to systematic risk. Both sets of estimates are based on market prices of equity and relate to end-2007. Finally, we abstract (for simplicity) from model and estimation uncertainty. Given these assumptions, we then derive the expected shortfall of the system and each institution's contribution to it. The results are shown in Graph 3, which plots each institution's contribution to system-wide risk against three of its drivers, namely the institution's size, probability of default and exposure to the common factor.

... illustrates that there is no single proxy for systemic importance

The results indicate quite clearly that the interaction of the various factors plays a key role. None of them, in isolation, provides a fully satisfactory proxy for systemic importance. For example, the largest institution in the system illustrated in Graph 3 is also the one with the biggest contribution (red dot). However, owing to its comparatively high probability of default, the institution with the fourth largest contribution is also one of the smallest and the least exposed to the common risk factor (blue dot). This highlights an important strength of the Shapley value methodology, namely that it allows for a straightforward *quantification of the interactions* of the various drivers.

Conclusion

This paper has presented a very general methodology to quantify the contribution of individual institutions to systemic risk. For a given measure of systemic risk, this is equivalent to calculating their systemic importance. The methodology can be applied to a wide variety of measures of systemic risk, and is very intuitive and flexible. As shown elsewhere, it subsumes other much more restrictive procedures as special cases (Tarashev et al (2009)). The methodology is very helpful in structuring an analysis of the drivers of systemic importance and in quantifying their relative impact.

In practice, any measure of individual institutions' systemic importance will necessarily be based on a specific measure (or measures) of systemic risk. The construction of such measures faces a number of tough challenges. These largely reflect the need to define what the relevant system is and to estimate the appropriate parameters. In the specific setting used here, these parameters include the probability of default and loss-given-default of individual institutions, exposures to common risk factors and the size distribution of the system. We have discussed how some of these challenges can be met and illustrated this with a concrete but simplified example using real-world data. In future, tools such as this one will inevitably be part of the arsenal of weapons needed to implement a financial policy framework with a macroprudential orientation, as called for by the international policy community.

Technical appendix: expected shortfall

Expected shortfall, also known as expected tail loss, is the measure of systemic risk we use in all numerical examples. It is defined as the expectation of default-related losses in the system, conditional on a systemic event. This event occurs when system-wide losses equal or exceed some (in this article, the 98th) percentile of their probability distribution.

We specify this probability distribution as follows. System-wide losses equal $\sum_{i=1}^{N} s_i \cdot LGD_i \cdot I_i$, where s_i is the size of the liabilities of institution *i*, LGD_i (loss-given-default) is the share of s_i that is lost if that institution defaults, and I_i is an indicator variable that equals 1 if institution *i* defaults and 0 otherwise. Without loss of generality, the overall size of the system is set to unity, $\sum_{i=1}^{N} s_i = 1$, and, for simplicity, it is assumed that $LGD_i = 55\%$ for all institutions. Finally, in line with structural credit risk models, institution *i* is assumed to default when its assets V_i fall below a particular threshold. Specifically, this happens when $V_i = \rho_i \cdot M + \sqrt{1-\rho_i^2} Z_i < \Phi^{-1}(PD_i)$, where the value of assets is driven by one risk factor that is common to all institutions, M, and another risk factor that is specific to institution *i*, Z_i, and both factors are standard normal variables. In addition, PD_i denotes the unconditional probability of default of institution *i* and Φ^{-1} is the inverse of the standard normal CDF. Finally, the loadings on the common (or systematic) factor, $\rho_i \in [0,1]$ for $i \in \{1, \cdots, N\}$, determine the correlation of defaults within the system.

We quantify expected shortfall using Monte Carlo simulations that take as inputs the following parameters for each institution *i*: s_i, LGD_i, PD_i, ρ_i.

References

Acharya, V and M Richardson, eds (2009): *Restoring financial stability: how to repair a failed system*, John Wiley & Sons.

Borio, C (2009): "Implementing the macroprudential approach to financial regulation and supervision", *Banque de France Financial Stability Review*, September.

Borio, C and M Drehmann (2009): "Towards an operational framework for financial stability: 'fuzzy' measurement and its consequences", *BIS Working Papers*, no 284, June.

Crockett, A (2000): "Marrying the micro- and macro-prudential dimensions of financial stability", *BIS Speeches*, 21 September.

De Larosière, J (2009): "The high-level group report on financial supervision in the EU".

Group of 20 (2009): "G20 working group 1: enhancing sound regulation and strengthening transparency", 25 March.

Knight, M (2006): "Marrying the micro- and macroprudential dimensions of financial stability: six years on", *BIS Speeches*, 5 October.

Kuritzkes, A, T Schuermann and S Weiner (2005): "Deposit insurance and risk management of the US banking system: what is the loss distribution faced by the FDIC?", *Journal of Financial Services Research*, no 27:3, pp 217–42.

Mas-Colell, A, A Whinston and J Green (1995): *Microeconomic theory*, Oxford University Press.

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 (2009): "Measuring portfolio credit risk correctly: why parameter uncertainty matters", *BIS Working Papers*, no 280.

Tarashev, N, C Borio and K Tsatsaronis (2009): "Allocating systemic risk to individual institutions: methodology and policy applications", *BIS Working Papers*, forthcoming.