

Attributing Systemic Risk to Individual Institutions¹

Methodology and Policy Applications

Nikola Tarashev, Claudio Borio and Kostas Tsatsaronis²

This draft: February 12, 2010

Abstract

Tools that attribute system-wide risk to individual institutions are key elements of an operational macroprudential approach to financial stability. We propose to measure institutions' systemic importance via an attribution methodology that is based on a game theoretic construct: the Shapley value. This methodology has a number of appealing features. First, owing to weak underlying assumptions, it can be used in conjunction with all popular risk measures. Second, it provides an exact allocation of risk that satisfies a concrete fairness criterion. Third, it accommodates easily uncertainty about model parameters and, more generally, uncertainty about the validity of the risk model.

We apply this methodology to a number of stylised banking systems and two risk metrics, value-at-risk and expected shortfall. This allows us to study the interaction of different drivers of systemic importance: size, risk profile and common-risk-factor exposures of individual institutions. We also prove a theorem that, under mild conditions, the systemic importance of an institution increases more than proportionately with its size. An implication of this finding is that a policy intervention calibrated to be proportional to the size of an institution would understate the systemic importance of large firms relative to that of small ones. In addition, we demonstrate how the Shapley value attribution methodology can be used for the calibration of macroprudential capital requirements that have an objective defined at the level of the overall system but are applied at the level of individual institutions. Finally, we argue that different policy tools call for using different applications of the general Shapley value methodology.

Keywords: Systemic importance, macroprudential approach, Shapley value

JEL Classification: C15, C71, G20, G28.

¹ The authors thank Marek Hlavacek for outstanding research assistance and Mathias Drehmann, Andrew Lo, Goetz von Peter and seminar participants at the AEA meetings in Atlanta, BIS, ECB and the SUERF colloquium in Utrecht for helpful comments on earlier drafts. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank for International Settlements.

² Centralbahnplatz 2, Bank for International Settlements, CH-4002 Basel, Switzerland.

E-mail: nikola.tarashev@bis.org, claudio.borio@bis.org, ktsatsaronis@bis.org.

Introduction

A key policy lesson from the recent financial crisis has been the need to put greater emphasis on a systemic approach to financial stability. The chain of events that started with problems with portfolios of sub-prime mortgages quickly developed into a systemic crisis that engulfed financial institutions and markets across the world, triggering a severe recession. Building better defences against systemic risk has emerged as a policy priority and the objective of strengthening the macroprudential orientation of financial stability frameworks has risen to the top of the international policy agenda.³

An operational macroprudential policy framework requires a gauge of the systemic importance of individual institutions. The reason is that key aspects of the *instruments* available to policymakers are determined at the firm level. This is true of tools to mitigate *ex ante* the risk of systemic disruptions, such as regulatory minimum capital and liquidity requirements, and of *ex post* supervisory interventions to contain the systemic externalities from distress in specific institutions. The decision of US authorities to take the unprecedented step of offering emergency financial support to AIG provides a case in point. The decision was motivated by concerns about the repercussions of the failure of this institution on its extensive list of counterparties in the credit derivatives market. In other words, it was a concern about the systemic importance of the institution that guided the intervention.

Measuring systemic importance by attributing system-wide risk to individual institutions is akin to a problem already tackled by game theorists. In his search of a solution to cooperative games, Lloyd Shapley (1953) developed an attribution methodology that carries his name: Shapley value. The portion of the overall value (e.g. output) that this methodology attributes to each player in a game equals the average of this player's marginal contributions to the value created by all possible subsets of players. This results in a fair allocation of value in the sense that the value created jointly by two players is split equally between them.

In order to measure individual institutions' systemic importance, this paper transposes the Shapley value methodology to the field of risk attribution. In addition to its fairness property – whereby the risk created jointly by two institutions is split equally between them – the methodology possesses a number of other desirable features. It is simple, yet efficient in the

³ See BIS (2009), G20 (2009), De Larosiere (2009), FSB (2009). 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. See Crockett (2000), Knight (2006), and Borio (2003 and 2009) for an elaboration of the macroprudential approach and progress in its implementation.

sense that the shares of systemic risk attributed to individual institutions add up *exactly* to the total. It is flexible since the sufficient conditions for its application are so weak that it can be applied to any measure of risk that treats the system as a portfolio of institutions. It also encompasses all attribution procedures that have been studied in different contexts in the literature. Finally, it can deal with model and parameter uncertainty as it can easily combine information from different risk models and address estimation noise in order to produce robust assessments of systemic importance.

Besides introducing the Shapley value to the field of systemic risk, the paper makes three main contributions.⁴ The first contribution relates to the analysis of different drivers of systemic importance. We apply the Shapley value methodology in a number of stylised settings in order to highlight the role that an institution's size, risk profile and strength of exposure to a common risk factor play in shaping the institution's contribution to system-wide risk. Quite intuitively, greater size, probability of default (PD) and exposure to systematic risk raise the systemic importance of an institution, with the impact of one driver being reinforced by that of others.⁵ A more subtle finding of the analysis is a non-linear – *convex* – relationship between institution size and contribution to system-wide risk. This is a general result that is derived in the form of a theorem, drawing on the Shapley methodology. It implies that a policy intervention calibrated to be proportional to the size of an institution would understate the systemic importance of large firms relative to that of small ones. From a methodological perspective, the result also demonstrates the usefulness of the Shapley value. By casting the attribution problem in terms of a set of marginal contributions, the Shapley value makes it possible to analyse the impact of individual drivers of risk in a tractable way.

A second contribution of the paper is to illustrate, in a stylised setting, the implications of different policy tools that are based on minimum capital requirements and target financial stability. The three tools we consider are applied at the level of individual institutions and share one objective: a particular level of risk at the level of the overall system. The first tool attains this objective while equalising the riskiness of individual institutions. The second tool conditions on the same level of systemic risk but equalises the systemic importance of

⁴ Strictly speaking the Shapley methodology can also be used to attribute any portfolio risk measure to the components that constitute the portfolio. We do not pursue these extensions because their analysis would require the examination of the implications of the specific institutional context in which these attributions are performed.

⁵ Throughout the paper there is a distinction between the terms systemic risk and systematic risk. The former refers to the risk that problems will arise that will impede the ability of the financial system to function. The latter refers to the commonality in risk exposures of financial institutions (in the same spirit as the "market" is analysed in the CAPM). This means that systemic risk can have systematic and idiosyncratic components.

individual institutions (controlling for their size). Finally, the third tool minimises aggregate capital holdings, given the target level of systemic risk. An interesting result is that, when institutions differ only with respect to their exposures to a common risk factor, the capital charges that equalise institutions' systemic importance are: (i) associated with a lower level of aggregate capital than the charges that equalise individual riskiness; (ii) quite close to the charges that explicitly target a minimum level of aggregate capital.

As a third contribution, the paper analyses, within a common framework, two alternative attribution procedures. The two procedures, which have been studied separately in the literature, are seen to be special applications of the Shapley value methodology. We show that one of the procedures captures the *contribution* of individual institutions to systemic risk, whereas the other one reflects institutions' *participation* in systemic events. In gauging systemic importance, the first procedure combines the risk that an institution generates on its own with the incremental risk generated by this institution in any possible subset of the system. The procedure thus captures the *impact* of the institution on system-wide risk, i.e. on the likelihood and severity of systemic events. It is, therefore, suited for the calibration of macroprudential tools that are designed to limit this impact. By contrast, the second procedure calculates the expected contribution of an institution to the overall cost of systemic events, taking such events as given. This makes it the procedure to use in deriving *actuarially fair premia* for insurance against systemic events.

The objective of the paper is not to propose a measure of systemic risk but to present a methodology of attributing this risk, however it is measured, to individual institutions. For the purposes of our numerical analysis and only as an illustration, we use a specific model of system-wide losses and specific metrics that we apply within this model in order to gauge systemic risk. The metrics we choose – value-at-risk (VaR) and expected shortfall (ES) – essentially measure risk as the (expected) loss on the aggregate exposure to the institutions in a system, conditional on certain tail events. We argue that ES is an intuitively appealing approach to measuring systemic risk but we also analyse VaR as an alternative. Most results of the analysis do not depend on our choice of a model and risk metrics. A notable exception is the result on the convex relationship between size and systemic importance, which is independent of the chosen model but is derived only in the context of ES.

The rest of this paper is organized as follows. Section 1 reviews existing methods for the measurement of systemic risk and the attribution of this risk to individual institutions. Section 2 develops a stylised model of systemic risk and then specifies two alternative metrics for this risk. The section also presents and studies three attribution procedures that deliver alternative measures of each institution's systemic importance. Section 3 and 4 analyse, respectively, how different aspects of the system affect its overall risk and the systemic

importance of individual institutions. Finally, Section 5 provides examples of how the attribution of systemic risk can be used in prudential policy tools.

1. Related literature

The related literature can be divided into two streams. The first stream focuses on measuring total system-wide risk when the system is considered as a portfolio of institutions. The second stream studies procedures for attributing total system-wide risk to individual institutions. A key contribution of our paper is to propose a general attribution methodology that (i) can be applied to all of the systemic risk measures developed in the first stream of the literature and (ii) subsumes as special cases all previously studied attribution procedures.

Measuring overall risk: from investment portfolios to financial systems

The literature has developed several measures of systemic risk. Of particular interest are those that treat explicitly the financial system as a portfolio of institutions. Examples include the measures used in Geluk et al (2009), Kuritzkes et al (2005), BIS (2008, 2009), Goodhart and Segoviano (2008), and IMF (2008, 2009). In the context of the methodology developed in this paper, these measures of systemic risk are relevant for two reasons. First, they all provide a *single* metric of systemic risk that encompasses *all* institutions in the system. Second, they can be applied to *any* subset of institutions in the system. Given these two features, the quantum or risk implied by a given measure can be allocated across institutions on the basis of the Shapley value methodology.

Attributing risk

An attribution method decomposes the aggregate quantum of risk in order to allocate it across individual contributors. Even though a number of such methods have been discussed in the literature, they have been applied mostly in the context of investment portfolios. As pointed out by Acharya and Richardson (2009), however, the close correspondence between measures of portfolio risk and measures of systemic risk leads naturally to a correspondence between the respective attribution methods. In this section we discuss attribution methods that have been applied to either of the two types of risk measures.

The most popular method for allocating risk across individual investment exposures considers the losses each one of them is expected to generate in an event of general distress (Praschnik et al (2001), Hallerbach (2002), Kurth and Tasche (2003) and Glasserman (2005)). The method has been recently used by Acharya et al (2009) to obtain indirect measures of the systemic importance of financial institutions. It is also used by

Huang et al (2009) in the context of Asia-Pacific banks. An appealing feature of this method is that the portions of risk it attributes to different exposures add up *exactly* to the chosen measure of portfolio risk. A disadvantage is that the method cannot be applied to cases where system-wide risk is not measured by reference to a fixed set of events. This would be the case when the choice of risk metric is the variance or higher moments of the portfolio (system) loss distribution. We show below that this attribution method is a specific application of the Shapley value methodology.

Koyluoglu and Stoker (2002) decompose the variance of losses on an investment portfolio using several approaches, one of which is based on the Shapley value. This, alternative, application of the Shapley value averages the contributions of an exposure to the variance of the losses on all sub-portfolios to which this exposure belongs. A key difference between Koyluoglu and Stoker (2002) and this paper is that they do not focus on tail risk and thus do not condition the risk measure on extreme losses or situations of generalised distress.

Another decomposition method has been proposed by Gordy and Lütkebohmert (2007). They make use of the asymptotic single risk factor (ASRF) model and a so-called "granularity adjustment" (GA). In addition to incorporating a single common risk factor, the ASRF model hinges on the assumption that the portfolio is perfectly granular, in the sense that there is a large number of exposures and the size of the largest exposure is vanishingly small (Gordy (2003)). When the measure of systemic risk is VaR, the GA provides an *approximate* correction for the inaccuracies that arise from violations of the perfect-granularity assumption. Developed in the context of portfolio risk, the ASRF-GA method has not been previously considered for the attribution of systemic risk. We analyse this method as an approximation to a specific application of the Shapley value methodology and, in line with Martin and Wilde (2002), we find that it works well when the violation of the perfect-granularity assumption is not too strong.

A rather different approach underpins CoVaR, which has been applied by Adrian and Brunnermeier (2008) to the market risk of an investment portfolio and suggested as a way to measure the systemic importance of institutions. Applied to a financial system, CoVaR would gauge the severity of distress in one institution, conditional on distress in another institution or in a group of institutions. For example, a CoVaR measure could equal the VaR of losses in bank A conditional on the losses in the entire banking system being equal to their VaR level. Since CoVaR captures the tail interdependence between losses on bank A and those on the banking system, it is a specific measure of the systemic importance of bank A.

The approach embedded in CoVaR and the one we take in this paper are fundamentally different. In this paper, we adopt a top-down approach that gauges systemic importance by

attributing system-wide risk to individual institutions. By contrast, CoVaR focuses directly on individual institutions (or groups of institutions) and is not based on a specific measure of system-wide risk. It is a bottom-up approach that does *not* deliver components that add up to the total. In terms of the above example, adding the CoVaRs of all the banks in a system will not deliver the system-wide VaR.

2. Systemic risk and systemic importance

This section lays out the analytic foundations of the analysis. The first subsection defines two popular measures of risk, which the paper focuses on. The second subsection specifies the stochastic environment that drives the probability distribution of losses in the system. Then, the third subsection presents the Shapley value methodology as a tool for attributing systemic risk to individual institutions. The fourth subsection considers three concrete attribution procedures, two of which are particular applications of the Shapley value methodology.

2.1 Two concrete measures of systemic tail risk

Let a financial system be populated by n institutions (henceforth, “banks”), indexed by $i \in \{1, 2, \dots, n\}$, and incur losses only when one or several of these banks default. The loss associated with bank i equals

$$L_i = s_i \cdot LGD_i \cdot I_i, \tag{1}$$

where s_i stands for the size of the liabilities of bank i , LGD_i is the share of bank i liabilities lost if it defaults, and I_i is an indicator variable that equals unity when bank i is in default and zero otherwise.

A measure of systemic risk should incorporate the joint probability distribution of losses, $\{L_1, L_2, \dots, L_n\}$. As stressed in Section 2.3 below, the Shapley value methodology can be applied to *any* such measure as long as it is defined on each subset of $\{L_1, L_2, \dots, L_n\}$.

In this paper, we derive numerical results for two popular measures of tail risk: value-at-risk (VaR) and expected shortfall (ES). Each of these measures is defined by a different set of tail events. VaR at confidence level q^{VaR} equals the level of losses that is exceeded with probability $(1 - q^{VaR})$. Thus, the tail events under the VaR measure are those associated with the q^{VaR} quantile of the probability distribution of losses. For the numerical exercises below, we assume that $q^{VaR} = 0.999$. In turn, ES is the expectation of losses, conditional on them being above the q^{ES} quantile of their distribution. Thus, a tail event under the ES measure

materialises if and only if losses exceed this quantile. For the numerical exercises below, we assume that $q^{ES}=0.998$.⁶ When either of the two measures is applied to the overall system, the underlying tail events will be referred to as “systemic events”.

This paper does not take a stand on whether VaR or ES is the appropriate measure of systemic tail risk. Being focused on a specific quantile, VaR reveals the smallest loss in the tail of the loss distribution but provides no information about the overall severity of the losses in this tail. This issue is addressed by ES, which yields a summary statistic (the mean) of loss severity in the tail.⁷ However, an important drawback of ES is that it is estimated with substantial noise in real-world applications that rely on actual data of losses. This drawback is substantially smaller in the case of the VaR, precisely because its estimation is that of a quantile, as opposed to a mean (Heyde et al (2006)).

2.2 Towards a probability distribution of systemic losses

We apply the VaR and ES measures to a probability distribution of systemic losses, which we define on the basis of the following stochastic environment. In line with the tradition of structural credit risk models, we assume that bank i defaults if and only if its assets V_i fall below the default point DP_i . Specifically:

$$I_i = 1 \text{ if and only if } V_i < DP_i \text{ and } I_i = 0 \text{ otherwise} \quad (2)$$

In addition, it will be assumed that V_i is driven by one risk factor that is common to all banks, M , and another risk factor that is specific to bank i , Z_i . Concretely:

$$V_i = \rho_i \cdot M + \sqrt{1 - \rho_i^2} Z_i, \text{ for all } i \in \{1, 2, \dots, n\} \quad (3)$$

where each risk factor is a standard normal variable and all factors are mutually independent.⁸ The common-factor loadings (or exposures), $\rho_i \in [0, 1]$ for all $i \in \{1, 2, \dots, n\}$, imply that the asset correlation for any two banks i and j equals $\rho_i \cdot \rho_j$. Common-factor exposures, which explain how shocks external to the system can systematically give rise to joint failures, parallel a key building block of portfolio credit risk models.

⁶ The adopted difference between the two quantiles q^{VaR} and q^{ES} renders the values of VaR and ES measures comparable. None of the conclusions in this article hinges on the relative values of q^{VaR} and q^{ES} .

⁷ A related issue that the so-called “sub-additivity” property is violated by VaR but not by ES (see Hull (2006)).

⁸ This assumption circumvents important empirical questions related to the *shape* of probability distributions of asset returns and the associated uncertainty (see, for example, Hull and White (2004) and Tarashev and Zhu (2008)). As discussed below, however, such uncertainty can be incorporated in the Shapley value methodology that is at the heart of the paper.

We acknowledge that such a setup is likely to miss an important feature of financial systems that distinguishes them from investment portfolios. Concretely, banks may be related not only via their exposure to common risk factors that are external to the system but also via interbank exposures, which propagate shocks within the system and create so-called domino effects. Interbank exposures, which imply that the financial system should be considered not only as a portfolio but also as a *network* of intuitions,⁹ are likely to have a material impact on the level of systemic risk and on the systemic importance of individual institutions. We abstract from this impact in order to illustrate the Shapley value methodology in a parsimonious setting.

Expressions (1)-(3) define the joint probability distribution of losses, $\{L_1, L_2, \dots, L_n\}$. Two additional assumptions limit the computation burden without influencing the main messages of the analysis. First, loss-given-default is set to $LGD_i = 55\%$ for all i . Second, the overall size of the system is normalised to unity, $\sum_{i=1}^n s_i = 1$, without loss of generality.

The inputs required for the calculation of any the above measures of systemic risk are the size of each institution, its probability of default, the loss given default in each case, and an estimate of the likelihood of joint defaults. The likelihood of joint defaults is typically derived from the correlation of banks' asset returns, which can be estimated from equity and debt prices (as done, for example, by Moody's KMV in their GCorr model). This practice, however, may change in the future, given evidence from the current crisis that, at a time of stress, the degree of interconnectedness in the banking system is largely determined by features of the liability side of balance sheets. This issue notwithstanding, any specific data that are relevant for the estimation of default correlations may be complemented with information from supervisory assessments.

2.3 The Shapley value approach: a general attribution procedure

The Shapley value methodology was developed in the context of cooperative games, in which the collective effort of a group of players generates a shared "value" (e.g. wealth) for the group as a whole.¹⁰ Given such a value, the methodology decomposes it in order to allocate it across players according to their individual contributions. The share of the aggregate value attributed to a particular player is this player's Shapley value.

⁹ For an in-depth analysis of the network structure of a national interbank market, see Boss et al (2004).

¹⁰ The discussion of Shapley value in this paper draws heavily on Mas-Colell et al (1995), pages 679-684. The Shapley value was first introduced in Shapley (1953).

The Shapley value methodology can be applied directly to a financial system. In this context, the players are institutions which engage in interrelated risky activities that drive systemic risk. Then, in the light of Section 2.1, the “value” of this risk is system-wide VaR or ES. Finally, the systemic importance of each institution is its Shapley-value.

This subsection first outlines the Shapley value methodology, stating explicitly the limited sufficient conditions for its applicability and listing its properties, which carry much intuitive appeal. Then, the section turns to the fact that the generality of the methodology makes it possible to decompose a given system-wide VaR or ES in different ways. The section concludes by arguing that the applicability of different decompositions – and, thus, different measures of systemic importance – depends on the problem at hand.

In order to apply the Shapley value methodology to a financial system, it is *sufficient* to define a so-called “characteristic function.” This function is the same for all possible subgroups of banks (or subsystems) and maps each subsystem into a risk measure. Given the setup developed since the beginning of Section 2, the characteristic function, ϱ , should accept as input any one of the 2^n subsystems of banks¹¹ and should deliver the system-wide VaR or ES when applied to the entire system. That said, it should be noted that ϱ could alternatively be based on *any* one of the existing measures of systemic risk presented in Section 1.1, simply because each one of them is defined for any subgroup of institutions in a financial system.

The derivation of the Shapley values involves the following thought process. Suppose that banks are ordered at random and consider the subsystem S that comprises all the banks in front of bank i as well as bank i . The contribution of bank i to the risk of subsystem S equals the difference between the risk of subsystem S and the risk of this subsystem when bank i is excluded from it: $\varrho(S) - \varrho(S - \{i\})$. The Shapley value of bank i , henceforth ShV_i , equals the expected value of such a contribution when the $n!$ possible orderings occur with an equal probability.

In the special case of a system comprising three banks, the Shapley value of bank 1 equals:

$$ShV_1(\{1,2,3\}) = \frac{1}{6} \left(2 \cdot (\varrho(\{1\}) - 0) + (\varrho(\{2,1\}) - \varrho(\{2\})) \right. \\ \left. + (\varrho(\{3,1\}) - \varrho(\{3\})) + 2 \cdot (\varrho(\{2,3,1\}) - \varrho(\{2,3\})) \right)$$

where $1/n! = 1/6$ is the probability of each of the six possible orderings. The first difference in the last expression is associated with two orderings, $[1,2,3]$ and $[1,3,2]$. The second and third

¹¹ These subsystems are: $\emptyset, \{1\}, \{2\}, \{3\}, \dots, \{n\}, \{1,2\}, \{1,3\}, \dots, \{n-1,n\}, \dots, \{1,2,3,\dots,n\}$.

differences are associated with one ordering each: $[2,1,3]$ and $[3,1,2]$, respectively. Finally, the fourth difference is associated with two orderings, $[2,3,1]$ and $[3,2,1]$. It incorporates the fact that $g(\{2,3,1\}) = g(\{3,2,1\})$ or, more generally, that the value of the characteristic function does not depend on how banks are ordered in the subsystem (see the symmetry property below).

Most generally, the Shapley value – or the systemic importance – of any bank i equals:

$$ShV_i(\Sigma) = \frac{1}{n} \sum_{n_s=1}^n \frac{1}{c(n_s)} \sum_{\substack{S \supset i \\ |S|=n_s}} (g(S) - g(S - \{i\})) \quad (4)$$

where Σ denotes the entire financial system, $S \supset i$ are all the subsystems in Σ containing bank i , $|S|$ stands for the number of banks in subsystem S , and $c(n_s)$ is the number of subsystems comprising n_s banks.¹² In addition, the empty set carries no risk: $g(\emptyset) = 0$.

For a given characteristic function g , the Shapley values of individual banks are a *unique* set of measures of systemic importance. This set possesses the following properties:

1) *Additivity (or efficiency)*: The sum of Shapley values equals the aggregate measure of systemic risk: $\sum_{i=1}^n ShV_i(\Sigma) = g(\Sigma)$.

2) *Symmetry*: The labelling of banks does not matter. More precisely, if the characteristic functions g and \tilde{g} differ only in that the roles of banks i and h are permuted, then $ShV_i(\Sigma; g) = ShV_h(\Sigma; \tilde{g})$.

3) *“Dummy axiom”*: If a bank carries no risk, then its Shapley value is zero.

4) *Linearity of characteristic functions*: Suppose that initially there is a set of alternative characteristic functions, a linear combination of which delivers a new characteristic function. The new Shapley value of any bank equals the same linear combination of the Shapley values implied for this bank by the initial set of characteristic functions. For example, if $g = \alpha \cdot g_1 + \beta \cdot g_2$ and α and β are constants, then $ShV_i(\Sigma, g) = \alpha \cdot ShV_i(\Sigma, g_1) + \beta \cdot ShV_i(\Sigma, g_2)$ for any bank i .

The linearity property of the Shapley value methodology implies that measures of systemic importance can account in an internally consistent manner for the ubiquitous issue of model and parameter uncertainty. For instance, there is no clear evidence whether the vulnerability

¹² Concretely, $c(n_s) = \frac{(n-1)!}{(n-n_s)!(n_s-1)!}$.

of financial systems is associated mainly with institutions' assets (credit exposures) or liabilities (funding exposures). Likewise, there is no consensus whether shocks exogenous to the financial system or the propagation of shocks within the system are the primary drivers of systemic events. Given this, it becomes inherently difficult to pinpoint the statistical properties of these shocks and to restrict the estimation noise in the parameters of data generating processes. Ultimately, all these different sources of uncertainty would imply that a prudential authority may want to consider a range of alternative measures of systemic risk, i.e. a range of alternative characteristic functions. The linearity property of Shapley values would then allow the authority to incorporate all these characteristic functions in a single attribution procedure, with the associated weights, i.e. α , β in the above example, reflecting the authority's perception of the validity of any given function.

A different perspective on the Shapley value methodology reveals that it satisfies an intuitive fairness criterion. Namely, the decomposition is such that the portion of systemic risk caused by the simultaneous presence of any two institutions in the system is split equally between them. As illustrated in MasCollé et al (1995), a specific implication of this is that the increment of the Shapley value of institution i caused by the presence of institution k equals the increment of the Shapley value of institution k caused by the presence of institution i . Moreover, this is true even if the Shapley value is defined on any subgroup of institutions in the entire financial system Σ :

$$ShV_i(S) - ShV_i(S - \{k\}) = ShV_k(S) - ShV_k(S - \{i\}) \quad (5)$$

for all i and k ; and all $S \in \Sigma$, such that $i, k \in S$.

Besides its intuitive appeal, the property of Shapley values in expression (5) helps bring to the fore differences between alternative applications of the general Shapley value methodology. We develop this point in the next subsection.

2.4 Three ways to measure systemic importance

When the measure of systemic risk is VaR or ES, the Shapley values of individual institutions can be based on either of two different characteristic functions. The two characteristic functions coincide when applied to the entire system but differ, in terms of the underlying tail events, when applied to subgroups of institutions. The upshot is two different attribution procedures that decompose the same magnitude of systemic risk in different ways. We outline these two attribution procedures in turn. In order to alleviate the exposition, in this subsection, we discuss only the attribution of systemic VaR, keeping in mind that the ES case is conceptually equivalent. Then, we outline a third attribution procedure, which is an analytic approximation of one of the first two. Finally, we argue that the different measures of

systemic importance, delivered by the alternative attribution procedures, should be used in different settings.

In exploring each procedure, it is important to keep in mind that the underlying stochastic environment generating default losses (recall Section 2.2) simplifies considerably the derivation of Shapley values. Since it is assumed that each bank is subject only to shocks external to the system, the statistical properties of the losses associated with a given bank are *unaffected* by the other banks and, thus, stay constant across subsystems. This property of default losses would be foregone if the system were considered as a network of institutions. Since, in this case, banks would propagate shocks from/to other banks, the losses associated with a given bank would depend on which other banks are in the subsystem in focus.

Procedure 1: varying tail events

This procedure is underpinned by the characteristic function \mathcal{G}^1 , which is such that $\mathcal{G}^1(S) = VaR(S)$ for *any* possible subsystem S in Σ . It is important to note that \mathcal{G}^1 defines the tail events at the level of each subsystem and these events typically differ from the systemic events, ie the tail events at the level of the entire financial system. Procedure 1 has been employed by Koyluoglu and Stoker (2002) but in a different context (see Section 1.2 above).

A measure of systemic importance obtained under Procedure 1 reflects the *contribution* of individual banks to the severity of the systemic events. As implied by expression (4), Procedure 1 gauges the systemic importance of bank i by combining the VaR that bank i would generate on its own with the extra amount of losses at the q^{VAR} quantile that would be generated if bank i were to be added to any possible group of other banks in the system. As a result, the Shapley value of bank i under this procedure is a direct reflection of the extent to which this bank affects the VaR of the system.

To understand further the characteristic function \mathcal{G}^1 , it is useful to revisit the fairness property in expression (5). Owing to its treatment of tail events, \mathcal{G}^1 reflects the extent to which the joint presence of two banks i and k raises the risk in a subsystem. The Shapley value methodology then splits the incremental amount of risk equally between the two banks. Specifically, provided that the risk factors affecting banks i and k relate positively, $ShV_i(S; \mathcal{G}^1) - ShV_i(S - \{k\}; \mathcal{G}^1) = ShV_k(S; \mathcal{G}^1) - ShV_k(S - \{i\}; \mathcal{G}^1) \geq 0$ and the inequality is strict for a strictly positive number of subsystems $S \in \Sigma$, such that $i, k \in S$.

Procedure 2: fixed tail events

Procedure 2 is another application of the Shapley-value methodology, based on a different characteristic function, \mathcal{G}^2 . For any subsystem S , $\mathcal{G}^2(S)$ equals the expected loss in this

subsystem conditional on the tail events in the *entire* system Σ , ie conditional on the systemic events. It is the different treatment of tail events that drives the difference between characteristic functions \mathcal{G}^I and \mathcal{G}^{II} .

A measure of systemic importance obtained under Procedure 2 captures the degree to which a bank is expected to *participate* in the systemic events. To see why, note first that \mathcal{G}^{II} leads to a substantial simplification because $\mathcal{G}^{II}(S) - \mathcal{G}^{II}(S - \{i\}) = E(L_i \mid \text{systemic event})$, which depends on i but not on S . Then, by expression (4), the Shapley value of bank i is simply the loss it is expected to generate, conditional on the systemic events:

$$ShV_i(S; \mathcal{G}^{II}) = ShV_i(\Sigma; \mathcal{G}^{II}) = E(L_i \mid \text{systemic event}) \text{ for all } i \in S \text{ and all } S \in \Sigma.$$

The characteristic function \mathcal{G}^{II} underpins an application of the Shapley value methodology that satisfies the letter but not the spirit of the fairness property in expression (5). The fundamental reason is that, since it takes systemic events as given, \mathcal{G}^{II} cannot convey how bank k affects the contribution of bank i to these events and vice versa:

$$ShV_i(S; \mathcal{G}^{II}) - ShV_i(S - \{k\}; \mathcal{G}^{II}) = ShV_k(S; \mathcal{G}^{II}) - ShV_k(S - \{j\}; \mathcal{G}^{II}) = 0 \text{ for each } S \in \Sigma.$$

Indeed, this is a manifestation of the fairness property but an uninformative one.

A different issue, which arises only in the context of system-wide VaR, is that an application of Procedure 2 may give rise to non-trivial computational complications that necessitate approximations. The reason is that, if losses have a continuous probability distribution, the systemic events underpinning the VaR measure – i.e. those corresponding to the q^{VaR} quantile of the probability distribution of losses – are of zero probability. Therefore, expectations conditional on such events are impossible to derive exactly. Hallerbach (2002) shows that the problem can be tackled numerically via a procedure in which there is a trade-off between the accuracy and efficiency of the conditional expectation estimator.

Procedure 2 has been a popular tool for the attribution of the risk of investment portfolios to individual exposures and has been recently used by Acharya and Richardson (2009) and Huang et al (2009) in the context of systemic risk (see Section 1.2 above). However, previous derivations of the procedure – such as those in Praschnik et al (2001), Hallerbach (2002), Kurth and Tasche (2003) and Glasserman (2005) – have been based on the linearity of the expectations operator, not on the Shapley value methodology. By extension, the properties of Procedure 2 have not been analysed alongside those of Procedure 1. In Section 2.4.1 below, we compare the two procedures and argue that they should be used in different contexts.

Procedure 3: ASRF model with a granularity adjustment

This procedure, which does not make use of the Shapley value methodology and has been developed only for VaR measures, is an analytic approximation of Procedure 2. Under Procedure 3, the portion of system-wide VaR attributed to bank i equals $MVaR_i^{ASRF,GA} = MVaR_i^{ASRF} + MGA_i$. The first summand, $MVaR_i^{ASRF}$, is derived in Gordy (2003) in the context of the asymptotic single risk factor (ASRF) model and, thus, incorporates the assumption that the system is perfectly granular (or asymptotic). The second summand, MGA_i , is derived in Gordy and Lütkebohmert (2007) is an approximate correction for departures from this assumption, i.e. a “granularity adjustment”:

$$MVaR_i^{ASRF} = s_i \cdot LGD \cdot \Phi \left(\frac{\Phi^{-1}(PD_i) - \rho_i \Phi^{-1}(1 - q^{VaR})}{\sqrt{1 - \rho_i^2}} \right)$$

$$MGA_i = s_i^2 \cdot LGD \cdot f \left(\delta_i, \sum_{j=1}^n MVaR_j^{ASRF}, MVaR_i^{ASRF} - LGD \cdot PD_i \right)$$

where Φ stands for the standard normal CDF and the analytic function f and the parameters δ_i are defined in Gordy and Lütkebohmert (2007). Given that the system-wide VaR has been estimated, it is typically possible to find unique δ_i that preserve the internal consistency of the model and result in $\sum_{i=1}^n MVaR_i^{ASRF,GA} = VaR$.¹³

In the limit in which the granularity of the system is infinitely fine, and thus idiosyncratic risk is fully diversified away, the granularity adjustment declines to zero. In this limit, given that there is a single common risk factor, the ASRF model and attribution Procedure 2 coincide.¹⁴ Thus, Procedure 3 can be viewed as an approximation to Procedure 2. Section 2.4.2 below studies the accuracy of this approximation, which, to the best of our knowledge, has not been done before.

¹³ The parameters δ_i partially reconcile differences between the default generating process implied by the ASRF model and that implied by CreditRisk+, which is used for the granularity adjustment. In this paper, the parameters δ_i are calibrated so that there is a close match between the right tails of these distributions (see Gordy and Lütkebohmert (2007), equation (18)). Importantly, any possible calibration of δ_i introduces a conceptual issue. Namely, in line with their intended purpose to account for the degree of diversification in the system (or portfolio), these parameter depend on the common factor loadings. However, contrary to economic logic, they are also affected by individual PDs, the VaR confidence level and an additional ad hoc parameter.

¹⁴ The proof of Proposition 1 in Tarashev (2009) proves this claim as well.

2.4.1 Comparison between Procedures 1 and 2

This section illustrates differences between measures of systemic importance obtained under Procedures 1 and 2 and then analyses the reasons for these differences. The analysis is centred around the following two possible objectives of a prudential authority, the first one of which calls for the use of Procedure 1 and the second for the use of Procedure 2:

1. Attain a particular cross-sectional distribution of institutions' contributions to a given level of systemic risk. Section 5 below motivates such an objective from a macroprudential point of view.
2. Require banks to pay – at actuarially fair premia – for a scheme that insures against the losses in pre-specified systemic events. Being equal to the expected loss associated with a bank, conditional on the systemic events, the premium reflects the bank's participation in these events.

We consider the above two objectives in stylised examples that illustrate sharply that a bank's *contribution* to systemic risk (captured by Procedure 1) could differ substantially from its expected *participation* in the systemic events (captured by Procedure 2). The first such example is provided by Table 1, in which systemic risk is measured by VaR and, thus, the systemic events occur when system-wide losses equal the q^{VaR} quantile of their probability distribution. In this example, the system comprises 10 banks that differ only with respect to their size. These banks are divided into two groups of five and each of the banks in the first (second) group accounts for 7% (13%) of the total size of the system.

The left-hand panel of the table illustrates clearly that the two procedures can deliver quite different measures of systemic importance. In the considered system, which features relatively low default correlations, the systemic events correspond to the failure of two large banks (and a VaR of 14.3 cents on the dollar). Since these events exclude losses from small banks, applying Procedure 2 leads to the conclusion that these banks are of *no* systemic importance. The reason for this conclusion can be traced to the fact that Procedure 2 fails to convey the impact of a given bank on the risk generated by other banks (recall the discussion in Section 2.4). For the system at hand, Procedure 2 fails to convey that the level of systemic risk is partly the result of the *simultaneous* presence of the *two* groups of banks in the system. For example, this level would have *halved* if the group of small banks had been excluded from the system. Such instances of the impact of small banks on systemic risk would be captured only if the procedure considered tail events at the level of each subgroup of banks. Since this is what Procedure 1 does, it attributes positive systemic importance to small banks. Procedure 1 is then the natural choice under the first of the above objectives, which calls for measuring banks' contribution to systemic risk.

Comparison between Procedures 1 and 2: a VaR example				
All banks: PD = 0.27% and LGD = 55%				
Group A banks: $n_A = 5$; $s_A = 0.07$. Group B banks: $n_B = 5$; $s_B = 0.13$.				
	Low default correlation $\rho_A = \rho_B = 0.60$		High default correlation $\rho_A = \rho_B = 0.724$	
	Procedure 1	Procedure 2	Procedure 1	Procedure 2
Group A	34.34%	0.0%	28.15%	100%
Group B	65.66%	100%	71.85%	0.0%
<i>total VaR</i>	14.3 (100%)	14.3 (100%)	15.4 (100%)	15.4 (100%)
<p>Note: Each panel refers to a different banking system. Systemic risk is measured as total VaR at the 99.9% confidence level, in cents per dollar exposure to the system. The first two rows report the overall share of each group of banks in total VaR, as allocated by the procedure specified in the column heading. The number of banks in group j equals n_j, the size of a bank in group j is s_j and the exposure of a bank in group j to the common factor is denoted by ρ_j.</p>				

Table 1

That said, Procedure 2 is designed for the second of the above objectives, i.e. the calculation of actuarially fair insurance premia when the insurance is against losses incurred in systemic events. To see this, consider again the system in which correlation is low. Given that the systemic events occur when system-wide losses equal 14.3 cents on the dollar and big banks are the sole drivers of such losses, these banks should be the only ones to pay actuarially fair insurance premia.

The picture is symmetric when higher default correlations lead to a system-wide VaR (15.4 cents on the dollar) that corresponds to the losses from the failure of four small banks (see right-hand panel of Table 1). In this case, Procedure 2 implies that the systemic importance of big banks is nil. For the reasons discussed above, this outcome is simply another example of a mismatch between the expected losses generated by a bank in systemic events and the contribution of this bank to systemic risk. Again, the mismatch suggests that Procedure 1 should be used for the first of the above objectives, even though Procedure 2 is the one to use for the second objective.

It should be noted that allowing for stochastic LGD would alter the numerical results in Table 1. For example, it would dampen the distinction between the two groups of banks under Procedure 2. To see why, note that, if the probability distribution of LGD is continuous, losses from each bank will enter the set of systemic events underpinning the VaR at any confidence level. This would guarantee a strictly positive level of systemic importance for each bank under Procedure 2.

That said, two points should be kept in mind. As discussed in Section 1.2, a departure from a step-wise loss distribution (which would result from a continuous PDF of LGD) raises significant computational issues when Procedure 2 is applied to a VaR measure of systemic risk. Second, keeping such issues aside, stochastic LGD does not alter the fact that Procedure 2 is not designed to convey the degree to which the interaction among different banks raises systemic risk. Numerical results, available upon request, reveal that the differences between Procedures 1 and 2 illustrated in Table 1 are maintained in qualitative terms even for a stochastic LGD with substantial variance.

A second example illustrates sharply the fact that a bank’s contribution to system-wide ES is also not equal to the extent to which the bank is expected to participate in the corresponding systemic events (see Table 2). The 4 banks in the hypothetical system of this example differ with respect to their individual PDs and loadings on the common risk factor. In order to analyse differences between the two attribution procedures, it suffices to consider the bank with the highest and that with the lowest probability of default, dubbed C and D, respectively. Bank C also features the lowest exposure to the common factor, whereas bank D features the highest exposure.

Comparison between Procedures 1 and 2: an ES example				
All banks: $s = 0.25$ and LGD = 55%				
	Low risk system		High risk system	
	$PD_A = PD_B = 0.31\%$, $PD_C = 0.62\%$, $PD_D = 0.28\%$ $\rho_A = \rho_B = 0.65$, $\rho_C = 0.10$, $\rho_D = 0.74$		$PD_A = PD_B = 0.62\%$, $PD_C = 1.24\%$, $PD_D = 0.56\%$ $\rho_A = \rho_B = 0.65$, $\rho_C = 0.10$, $\rho_D = 0.74$	
	Procedure 1	Procedure 2	Procedure 1	Procedure 2
Banks A and B	53%	49%	54%	57%
Bank C	20%	26%	17%	12%
Bank D	27%	25%	29%	31%
<i>Total ES</i>	18.4 (100%)	18.4 (100%)	26.2 (100%)	26.2 (100%)
Note: Each panel refers to a different banking system. Systemic risk is measured as total ES at the 99.8% confidence level, in cents per dollar exposure to the system. The first three rows report the share of each bank (or group of banks) in total ES, as allocated by the procedure specified in the column heading. The size of a bank is denoted by s , the PD of bank j is PD_j and the exposure of bank j to the common factor is denoted by ρ_j .				

Table 2

When the general level of banks’ PDs is low, Procedure 1 attributes a *larger* share of systemic risk to bank D than to bank C (left-hand panel). The underlying reason is that, with

its greater dependence on the common risk factor, bank D is more likely to be part of joint failures than is bank C. This raises the contribution of bank D to systemic risk relative to that of bank C. For example, removing bank D from the overall system makes the ES drop from 18.4 to 15.3 cents on the dollar, while removing bank C induces a smaller drop, to 17.6 cents. Procedure 1 incorporates such facts directly by considering the extent to which each bank raises the ESs of various subsystems. This makes the procedure a natural choice in the context of the first of the above objectives, which calls for gauging individual contributions to systemic risk.

For the same system, Procedure 2 delivers a different conclusion: that the systemic importance of bank D is *smaller* than that of bank C. To see why, note first that the systemic events in the considered system correspond to losses generated by the failure of one or more banks. Then recall that the level of systemic importance under Procedure 2 equals the expected losses of each bank, conditional on the systemic events, but is independent of a bank's propensity to participate in these events with other banks. Given this, the high likelihood of *solo* failures by bank C in the systemic events drives its measured level of systemic importance above that of bank D. Nonetheless, the levels of systemic importance obtained under Procedure 2 do equal the actuarially fair premia that banks should pay to a provider of insurance against the systemic events (which relates to the second of the above objectives).

The distinction between Procedures 1 and 2 is less sharp if the banks in the system feature higher PDs and, as a result, the systemic events underpinning the system-wide ES is associated *only* with losses from the failure of two or more banks (right-hand panel of Table 2). In this case, Procedure 2 joins Procedure 1 in attributing a higher portion of systemic risk to the bank with a higher exposure to the common factor, ie bank D. The qualitative similarity between the two procedures notwithstanding, Procedure 1 points to a smaller difference between banks C and D. This is because, while Procedure 2 focuses on a bank's role in the ES of the overall system where only joint failures matter, Procedure 1 considers also subsystems where the level of ES is affected by losses from single failures. In comparison to the overall system, the contributions of banks C and D to the risk of such subsystems differ less because the two banks are assumed to be of equal sizes and to feature high PDs (concretely, $PD_C > 1 - q^{ES}$ and $PD_D > 1 - q^{ES}$).

2.4.2 Comparison between Procedures 2 and 3

As stated above, Procedure 3 approximates well Procedure 2 when the granularity of the financial system is sufficiently fine, ie when there is a large number of banks and all bank sizes are similar. The left-hand and centre panels of Table 3 illustrate that this condition is

met by a system of 24 banks that differ only with respect to their PDs but not quite by an analogous system of 10 banks. A similar conclusion (not illustrated in the table) is reached in the context of banking systems in which banks differ from each other only with respect to their exposure to the common risk factor. Importantly, when banks' relative sizes differ, the system may remain lumpy irrespective of the number of banks. In turn, this implies that Procedure 3 may approximate poorly Procedure 2 even for systems comprised of a large number of banks (Table 3, right-hand panel).

Comparison between procedures 2 and 3						
All banks: PD = 0.3%, LGD = 55%						
	$n_A = n_B = 5$ $s_A = s_B = 0.1$ $\rho_A = 0.5, \rho_B = 0.5$		$n_A = n_B = 12$ $s_A = s_B = 0.0417$ $\rho_A = 0.5, \rho_B = 0.7$		$n_A = n_B = 12$ $s_A = 0.0167, s_B = 0.0667$ $\rho_A = \rho_B = 0.6$	
	Procedure 2	Procedure 3	Procedure 2	Procedure 3	Procedure 2	Procedure 3
Banks in group A	39%	35%	33%	34%	5%	15%
Banks in group B	61%	65%	67%	66%	95%	85%
<i>Total VaR</i>	11 (100%)	11 (100%)	9.17 (100%)	9.17 (100%)	11 (100%)	11 (100%)
Note: Each panel refers to a different banking system. Systemic risk is measured as total VaR at the 99.9% confidence level, in cents per dollar exposure to the system. The first two rows report the overall share of each group of banks in total VaR, as allocated by the procedure specified in the column heading. The number of banks in group j equals n_j , the size of a bank in group j is s_j and the exposure of a bank in group j to the common factor is denoted by ρ_j .						

Table 3

3. Drivers of systemic tail risk

This section moves away from methodological considerations in order to analyse the ES of concrete, albeit highly stylised and hypothetical, banking systems. The section documents the impact of four different drivers of systemic tail risk, as measured by ES: banks' number, relative sizes, individual PDs and exposures to the common risk factor.¹⁵

¹⁵ An analysis of these drivers under the VaR measure yields similar insights. Importantly, the paper abstracts from a number of additional drivers of systemic risk, such as the relationship between the number of defaults and LGD and drivers stemming from the network structure of the financial system.

The properties of ES have been analysed at considerable length in the context of portfolio tail risk. Cast in the present context, one of these properties is that the level of systemic risk increases as the PDs of some or all of the banks rise. Another well-known feature is that higher exposure to common risk factors increases the likelihood of joint failures, which typically raises tail risk in the system and, thus, its ES. Further, greater lumpiness of the financial system – caused by a reduction in the number of banks or greater disparity of their relative sizes – raises tail risk by restricting diversification benefits.

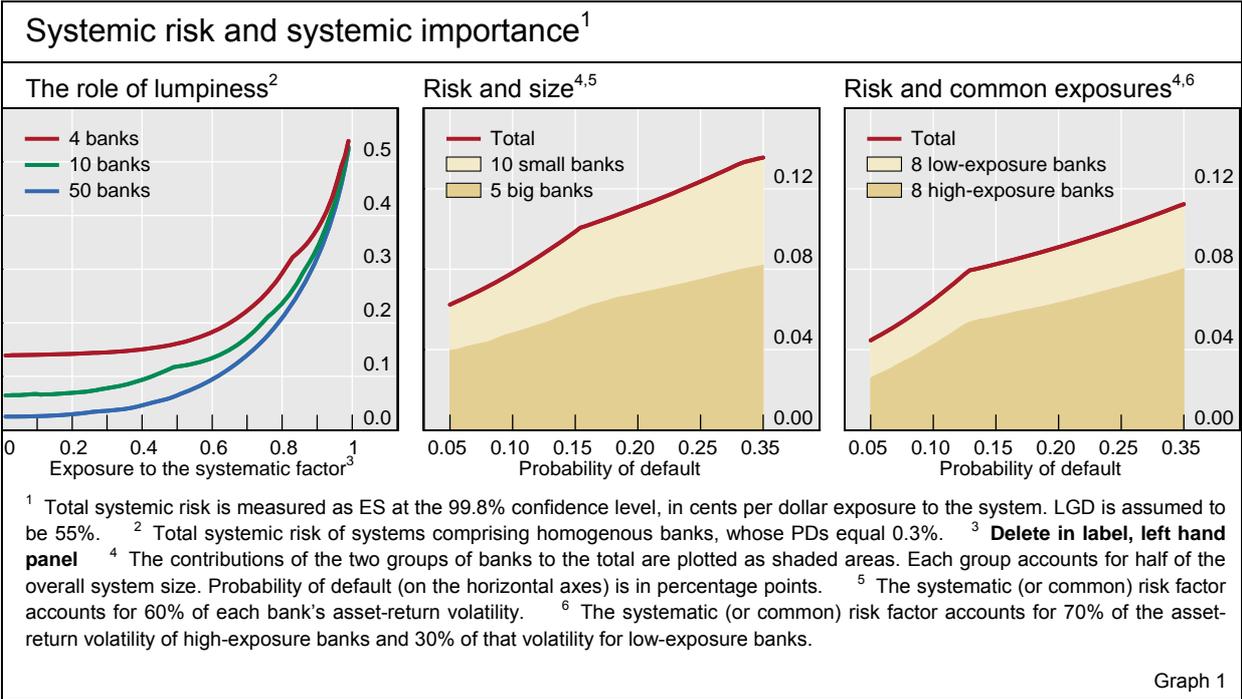
In order to illustrate additional properties of systemic risk (and, in the next sections, the attribution of systemic risk to individual banks), we resort to numerical examples that are based on specific values of banks' PDs and common-factor loadings. With the goal of staying in line with real-world bank characteristics, we calibrate hypothetical financial systems that are largely consistent with Moody's KMV estimates of the one-year PDs and asset-return correlations of 65 large internationally active banks at end-2007.¹⁶ These estimates suggest a typical (ie average) PD of 0.11% and a realistic high PD (ie average plus one standard deviation) of 0.3%. In addition, estimated asset-return correlations average 42% (consistent with a homogenous common factor loading, ρ , of $\sqrt{0.42} = 0.65$) and range between 14% ($\rho = 0.37$) and 55% ($\rho = 0.74$).

Benchmarking our calibration choices to these parameter estimates, we investigate the joint impact of system lumpiness and banks' exposure to the common factor on systemic tail risk. The results are portrayed in Graph 1, left-hand panel. In this panel, lumpiness is captured solely by the number of homogeneous banks in a hypothetical system and is held fixed (at one of three levels) in order to plot systemic risk as a function of the common-factor exposure.

A key message is that a decrease in the lumpiness of the system depresses systemic risk by more when banks' exposure to the common risk factor is smaller. In the limit case, in which all banks are exposed only to the common risk factor (i.e. when the asset-return correlations equal unity), changes in the lumpiness of the system are inconsequential. To see why, note that lower exposure to the common factor means greater importance of idiosyncratic risks. In turn, idiosyncratic risks are those that are diversified away at the level of the system when its lumpiness decreases (in this case, as the number of banks increases).

¹⁶ These estimates are delivered by the proprietary Credit Model and GCorr, respectively, and are based on market prices of banks' equity and debt.

The flipside of this intuitive result reveals an important insight regarding the consequences of measurement error. Namely, the different slopes of the three lines in the left-hand panel of Graph 1 indicate that systemic risk tends to increase faster in the exposure to the common factor when there are more banks in the system. Thus, a given error in the estimate of banks' exposures to the common factor is likely to result in a *larger* error in the measurement of systemic tail risk when the system is *less* lumpy.



4. Drivers of systemic importance

This section analyses drivers of systemic importance, measured here as the share of systemic ES attributed to individual banks by attribution Procedure 1.¹⁷ The four drivers considered below are those that were analysed in the context of systemic risk: i.e. banks' number, relative sizes, PDs and exposures to the common risk factor. The stylised banking systems that underpin the analysis are designed to meet two criteria. First, these banking systems are largely in line with Moody's KMV estimates of bank PDs and asset return correlations (see above). Second, the systems are populated by banks whose risk

¹⁷ Thus, in the light of the discussion in Section 2.2.1, systemic importance should be understood as being directly related to the institution's contribution to systemic risk.

characteristics are such as to allow for isolating the impact of specific drivers of systemic importance in a straightforward fashion.

4.1 The number of banks and their relative sizes

Quite intuitively, larger size implies greater systemic importance. We illustrate this in Table 4, for which we consider systems that possess the following three features. First, all banks in a given system share the same PD and exposure to the common factor. Second, there are 3 big banks of equal size, which account for 40% of the overall system. Third, a group of equally-sized small banks make up the rest of the system. In all of these systems, the systemic importance of a big bank is greater than that of a small one. More interestingly, as the number of small banks (but not their share in the overall size of the system) increases, their systemic importance declines both individually and as a group. The flipside of this is that the systemic importance of big banks rises.

System lumpiness										
Systemic risk and systemic importance										
	Low risk system (all banks: PD = 0.1%)					High risk system (all banks: PD = 0.3%)				
	$n_s = 5$	$n_s = 10$	$n_s = 15$	$n_s = 20$	$n_s = 25$	$n_s = 5$	$n_s = 10$	$n_s = 15$	$n_s = 20$	$n_s = 25$
3 big banks	43%	57%	63%	66%	68%	42%	52%	57%	59%	61%
n_s small banks	57%	43%	37%	34%	32%	58%	48%	43%	41%	39%
<i>Total ES</i>	9.8 (100%)	9.4 (100%)	9.3 (100%)	9.25 (100%)	9.23 (100%)	16.7 (100%)	15 (100%)	14.7 (100%)	14.4 (100%)	14.3 (100%)
<p>Note: Each column refers to a different banking system. Systemic risk is measured as total ES at the 99.8% confidence level, in cents per dollar exposure to the system. The first two rows report the overall share of each group of banks in total ES, as allocated by Procedure 1. The group of big banks accounts for 40% of the overall size of the system and the group of small banks accounts for 60%. Each bank is assumed to have the same sensitivity to the common risk factor, implying a common asset return correlation of 42% (or $\rho = 0.65$), and features an LGD of 55%.</p>										

Table 4

Further inspection of Table 4 reveals that the contribution to system-wide risk increases *more than proportionately* with relative size. To see this, consider the first column of the table, which relates to a system in which a big bank is 10% larger than a small one but is assigned a 23% greater share in systemic risk.¹⁸ This effect increases as banks' sizes become more

¹⁸ Concretely: $s_{big}/s_{small} = (0.4/3)/(0.6/5) = 1.11$ and $ShV_{big}/ShV_{small} = (43%/3)/(57%/5) = 1.25$.

disparate. In the fifth column of the table, which relates to a system where the sizes of big and small banks are roughly 5-to-1, the respective shares in systemic risk are 18-to-1.

The basic intuition for the relationship between size and systemic importance is that systemic (ie tail) events are associated with extreme losses, in which large banks are more likely to participate than smaller ones. This is an important property and a concrete example of how the macro-prudential perspective may provide unique insights that would be missed by a micro-oriented approach. If systemic importance increases faster than size, then prudential tools that aim at mitigating systemic risk should be designed so that their impact on an institution also increases more than proportionately with its size.

The convex, positive relationship between size and systemic importance is a robust result. It is supported by all the ES-related examples reported in the paper and by (unreported) extensive numerical investigation of the underlying risk model. Furthermore, we establish analytically that, when the metric of risk is ES, systemic importance increases at least proportionately with size under quite general conditions. To isolate the impact of size, we consider a general system and compare the relative contributions to system-wide risk of two banks that have identical risk profiles and differ from each other only in terms of their size. We then obtain the following result, which does *not* depend on specific assumptions about a number of drivers of systemic importance, such as the probability distribution of risk factors and the default correlation between institutions:

Theorem: *Let two banks differ only in terms of size. Suppose further that the contribution of either of these two banks to the ES of any other subgroup in the system decreases (weakly) as the number of banks in the subgroup increases. Then, the ratio of the Shapley value of the larger to that of the smaller bank is (weakly) bigger than the ratio of the respective sizes.*

The sufficient condition in the statement of the theorem is fairly weak and quite intuitive. In the appendix we show that it is a generalisation of the well-known sub-additivity of ES, or that the sum of the ESs of two portfolios is not smaller than the ES of a third portfolio that equals the sum of the first two.

The formal proof of the theorem, which is presented in the appendix, makes repeated use of the following fact. If the joint failure of the smaller bank with a group of other banks is a tail event, then the joint failure of the larger bank with the same group of other banks would also be a tail event. However, the converse need not be true. Or, as stated above, a larger bank appears in tail events more often than a smaller bank with an identical risk profile.

4.2 The exposure of banks to the common factor and their PDs

Another intuitive result is that systemic importance increases with the bank’s exposure to the common risk factor. This is illustrated in Table 5, in which each banking system is comprised of 20 banks, divided into two homogeneous groups, A and B, that differ only with respect to banks’ exposures to the common factor. Keeping the exposures to the common factor constant in group B but increasing them for group-A banks (across columns, in each panel) results in an increase in these banks’ share in systemic risk. In the specific example, their contribution rises from 44% to 60%.

Exposure to a common risk factor										
Systemic risk and systemic importance										
	Low risk system (all banks: PD = 0.1%)					High risk system (all banks: PD = 0.3%)				
	$\rho_A = 0.3$	$\rho_A = 0.4$	$\rho_A = 0.5$	$\rho_A = 0.6$	$\rho_A = 0.7$	$\rho_A = 0.3$	$\rho_A = 0.4$	$\rho_A = 0.5$	$\rho_A = 0.6$	$\rho_A = 0.7$
10 banks in group A	44%	46%	50%	54%	60%	42%	45%	50%	56%	63%
10 banks in group B	56%	54%	50%	46%	40%	58%	55%	50%	44%	37%
<i>Total ES</i>	4.0 (100%)	4.4 (100%)	5.0 (100%)	5.8 (100%)	6.8 (100%)	6.6 (100%)	7.2 (100%)	8.2 (100%)	9.8 (100%)	11.5 (100%)

Note: Each column refers to a different banking system. Systemic risk is measured as total ES at the 99.8% confidence level, in cents per dollar exposure to the system. The first two rows report the overall share of each group of banks in total ES, as captured by Procedure 1. The exposure of each of the 10 banks in group A to the single common risk factor is as given in the row headings. The exposure of each of the 10 banks in group B to the common risk factor is held fixed at $\rho_B = 0.5$. All banks are of equal size, $s = 0.05$, and feature LGDs of 55%.

Table 5

The reason for this result is straightforward. 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 means a higher likelihood of extreme losses, which leads to a higher level of systemic risk. Quite intuitively, the rise in the level of systemic risk is attributed mainly to the banks that are affected directly by the cause of this rise, ie those that experience an increase in their exposure to the common factor (ie group-A banks in Table 1).

Anticipating the analysis in the next section, it is important to also record that greater size or exposure to the common risk factor strengthens the positive impact of a higher PD on systemic importance. In order to illustrate how size and PD interact, Graph 1 (centre panel) considers a system in which banks differ only in terms of size. As PDs increase uniformly across all banks in this system, the portion of the expected shortfall attributable to larger banks increases by a bigger amount than that attributable to smaller banks. The right-hand panel of Figure 1 illustrates a similar point in the context of a system comprised of banks that

differ only with respect to their exposures to the common risk factor. Given that all of these banks experience the same rise in their PDs, the resulting increase in the contributions to systemic risk is greater for banks with a larger common-factor exposure.

5. Stylised policy tools

This section discusses how tools that attribute system-wide risk to individual institutions can be used in conducting prudential policy. More specifically, the section illustrates differences between micro- and macro-prudential approaches to achieving a specific level of risk by means of regulatory capital requirements. The basic premise is that by affecting institutions' risk profiles, capital charges affect the overall level of systemic risk and institutions' contribution to it.

We assume that the authorities apply capital charges to individual institutions with the objective of achieving a target for system-wide risk. We discuss three alternative approaches to calibrating these charges. The first approach equalises the risk at the level of each institution, ie attains the target level of systemic risk with a uniform PD across all institutions. We label this the “micro-prudential” approach in the sense that it is in the spirit of the current policy framework. The other two approaches take more of a “macro” perspective in attaining the same target for system-wide risk. One equalises the contributions of individual institutions to system-wide risk (ie equalises their Shapley values). The other minimises the overall level of capital in the system (ie makes sure that the marginal reduction of systemic risk by an extra unit of capital would be the same across institutions).

Our setup is parsimonious. We assume that there is a one-to-one mapping between the individual risk of a bank (its PD) and the amount of capital it holds. In addition, we assume that banks do not hold capital in excess of the level required by the authorities. Hence, changes in capital requirements have a direct effect on the leverage of banks and, ultimately, on their PDs. At the same time, we assume that capital requirements do not affect the size of balance sheets and banks' exposure to the common risk factor. Concretely, the mapping between a bank's capital and its probability of default, is given by:

$$PD_i = \Phi \left(\frac{\psi \cdot \left(1 - \frac{K_i}{V_i} \right) - 1}{\sigma_v} \right) \quad (6)$$

where V_i is the level of the bank's assets, σ_V stands for asset volatility, K_i is the level of equity capital and ψ is an adjustment factor.¹⁹ A policy intervention can alter K – and thus PD – at the level of each bank, but none of the other parameters. The analysis below is conducted with reference to systems populated by two groups of banks. Each of the groups accounts for half of the aggregate system assets and includes homogeneous banks. As detailed below, the two groups differ from each other in terms of specific risk parameters.

In the first example, the two sub-groups differ only in terms of the intensity of the exposure of the banks to the systematic risk factor. Banks in one group have a lower exposure to systematic risk (in terms of equation 3, $\rho = 0.30$), while banks in the other group have higher exposure to the systematic risk factor ($\rho = 0.70$). The two groups are identical to each other in terms of everything else.

The policy experiment is depicted on the left-hand side panel of Graph 2. The curve labelled *iso-ES* corresponds to the combinations of capital charges – those on banks with a low (high) exposure to the systematic risk factor are on the horizontal (vertical) axis — that achieve the target level of system-wide risk. The first policy approach attains the target level of systemic risk for equal PDs (i.e. for equal capital charges) at point A: the intersection between the *iso-ES* curve and the 45-degree line from the origin.

Equal capital charges, however, do not imply that banks contribute equally to systemic risk. The dashed curve labelled *equal ShV* denotes all capital allocations that result in equal contributions to system-wide risk by banks in the two groups (ie equal Shapley values).²⁰ Equalisation of Shapley values across different banks requires that banks with a greater exposure to the common factor (and hence higher contribution to system-wide risk) face higher capital charges that reduce their PDs (recall Graph 1).

The third approach to achieving the target level of system-wide risk is to seek an allocation that minimises the aggregate capital in the system. Graphically, this approach delivers the capital allocation given by the tangency point between the *iso-ES* curve and the straight line with a slope of -1 (ie the line perpendicular to the main diagonal) that is closest to the origin. This corresponds to point C.

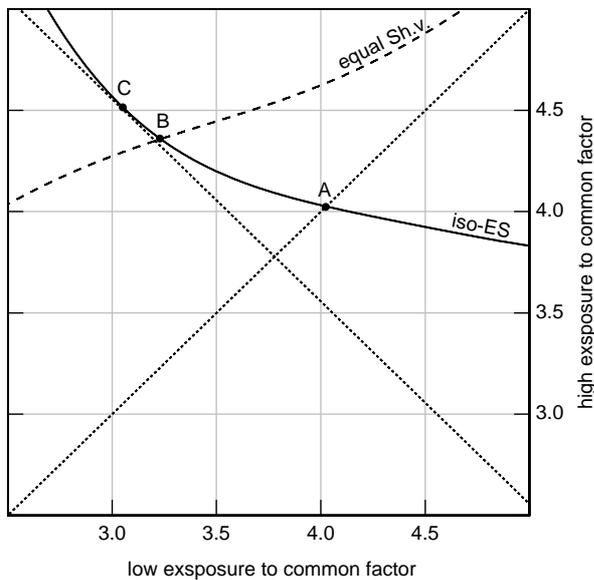
¹⁹ This equation is consistent with the model introduced in Section 2. Apparent differences stem from the fact that the formulae in Section 2 were designed to highlight how common-factor loadings enter the model, whereas here the emphasis is on the capital-to-asset ratio. To see the relationship between the alternative formulae, let the default point DP equal $\psi \cdot (V - K)$ and the asset volatility $1/V$ and note that the level of asset volatility is inconsequential for the analysis in previous sections. As initial conditions, we calibrate $K_i/V_i = 0.04$, $\sigma_V = 3.5\%$, and then set ψ so that $PD = 0.3\%$ for all banks.

²⁰ In the light of the discussion in Section 2.4.1, contributions to systemic risk are measured via Procedure 1.

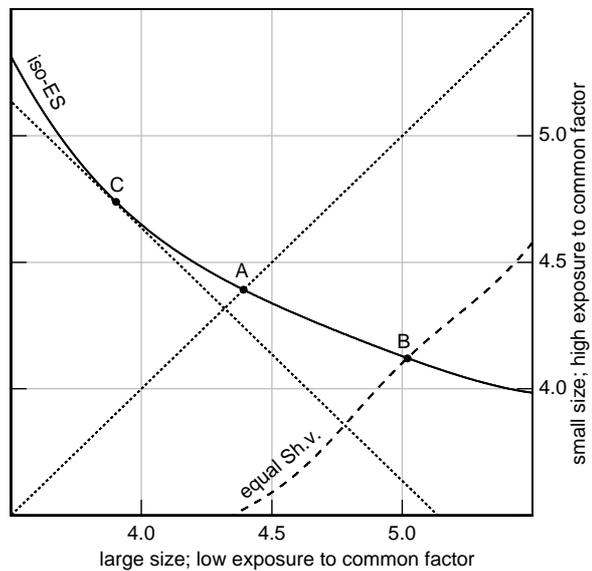
Macprudential policy interventions¹

Capital for two groups of banks

Banks differ in one aspect²



Banks differ in two aspects³



¹ Each panel corresponds to a specific system comprised of two groups of banks. Risk characteristics are identical across banks in each group. The aggregate group-wide assets are the same across groups. Each axis measures the capital charge for a particular group of banks, as per cent of the group's aggregate assets. The lines labelled "iso-ES" plot the pairs of capital charges, which imply that the system-level expected shortfall at the 99.8% confidence level equals 10% (left-hand panel) or 8% (right-hand panel) of aggregate liabilities in the system (LGD = 55% for all banks). The lines labelled "equal ShV" plot the capital pairs which imply that Shapley values are equal across the two groups of banks (equivalently, the ratios of Shapley values to corresponding bank sizes are equal across banks). ² The system comprises two groups of five banks each. The groups differ only with respect to the constituent banks' exposure to the common factor: $\rho_{low} = 0.30$ vs. $\rho_{high} = 0.70$. ³ The system comprises 20 banks. The first group comprises 4 large banks with a low exposure to the common factor: $s_{large} = 0.125$, $\rho_{low} = 0.30$. The second group comprises 16 small banks with a high exposure to the common factor: $s_{small} = 0.031$, $\rho_{high} = 0.70$.

Graph 2

In this example, both macroprudential approaches to achieving the target risk level (ie the second and third approaches) lead to efficiency gains in comparison to the microprudential approach (the first one). For the specific example used here, the aggregate level of capital in the system equals 4% of aggregate assets at point A, 3.8% at point B and 3.78% at point C (the minimum). The reason for this reduction in aggregate capital when going from A to B is related to the discussion in Section 4.2 about the interactions between size and PD in the determination of systemic importance. As illustrated in Graph 1 (right-hand panel), for a given change in PDs, banks that are more exposed to common risk factors experience a greater change in their contribution to systemic risk. Conversely, in equalising individual contributions to a fixed level of system-wide risk, the increase in the capital charge for banks with a greater common-factor exposure is smaller than the reduction for banks with a lower common-factor exposure. Hence, equalising contributions to systemic risk leads to a more efficient use of capital in this particular system than achieving the same risk with uniform capital levels.

The comparison between points B and C highlights the differences between the two macroprudential approaches. By construction, point C requires the lowest aggregate level of capital. By transferring more capital from the banks with a low systematic factor exposure to the group with a high exposure, this allocation exploits further the scope for efficiency than the one that equalises systemic risk contributions. That said, the two macroprudential approaches deliver the same ranking of the capital charges on the two types of banks: points B and C are on the same side of the 45-degree line.

The ranking of the three approaches in terms of aggregate capital requirements may differ from the above example if banks in the two groups differ in more than one aspects. For example a system where banks with a higher exposure to the common factor are also smaller would reverse the ordering of the first and second approaches in terms of efficiency in capital utilisation. This is shown in the right-hand panel of Graph 2. In this example, size and loading on the systematic factor have counteracting effects on the systemic importance of each bank. Hence, it is the relative importance of these effects that determines whether aggregate capital increases or declines when moving from the first, microprudential, approach to the second one, which equalises Shapley values. The specific parameterisation in this example leads to higher capital charges on the group of larger banks under the second approach (point B) and results in a higher level of aggregate capital (4.6% of aggregate assets) than those imposed by the first intervention (point A: 4.4% of assets). The third intervention that achieves the target with the minimum level of aggregate capital (4.3%) imposes higher capital on smaller banks because of their higher exposure to the systematic factor.

The policy examples in this subsection are intentionally cast in stylised settings that help highlight the interaction of different drivers of systemic importance. As such, the settings do not seek to capture particular empirical regularities and do not cover all possible ways in which institutions could respond to changing capital requirements. This leaves a number of important issues to future research.

Conclusion

Measures of the systemic importance of financial institutions are key inputs to macroprudential policy instruments. This paper proposes a general and flexible methodology for obtaining such measures by attributing systemic risk to individual institutions. The paper also demonstrates that different applications of the attribution methodology adopt different notions of systemic importance and, as a result, should be used for different macroprudential objectives. In addition, numerical examples highlight the importance of policy rules and

interventions that reflect not only the probability of a failure by an individual institution but also its exposure to common risk factors. The analysis also suggests that charges imposed on large institutions would need to reflect the more than proportionate impact of their size on systemic risk.

References

- Acharya V, L Pedersen, T Philippon, and M Richardson. 2009. "Measuring Systemic Risk", mimeo, New York University
- Acharya V and M Richardson, editors. 2009. Restoring Financial Stability: How to Repair a Failed System. John Wiley & Sons.
- Adrian T and M Brunnermeier. 2008. "CoVaR". Mimeo.
- Bank for International Settlements. 2008. Annual Report.
- Bank for International Settlements. 2009. Annual Report.
- Borio C. 2003. "Towards a macroprudential framework for financial supervision and regulation?". CESifo Economic Studies. vol 49, no 2/2003. pp 181-216. Also available as BIS Working Papers, no 128. February.
- Borio C (2009): "Implementing the macroprudential approach to financial regulation and supervision", *Banque de France Financial Stability Review*, September.
- Boss M, H Elsinger, M Summer and S Thurner (2004): "An Empirical Analysis of the Network Structure of the Austrian Interbank Market", Oesterreichische Nationalbank, Financial Stability Report 7.
- Crockett A. 2000. "Marrying the micro- and macroprudential dimensions of financial stability." BIS Speeches. 21 September.
- Financial Stability Board. 2009. "Report of the Financial Stability Forum on addressing procyclicality in the financial system." Basel. April.
- G20. 2009. "G20 working group 1: enhancing sound regulation and strengthening transparency." 25 March.
- Geluk J, L de Haan and C de Vries. 2009. "Weak and Strong Systemic Fragility". Mimeo.
- Goodhart C and M Segoviano. 2008. "Banking Stability Measures". IMF Working Paper.
- Gordy, M. 2003. "A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules." *Journal of Financial Intermediation* 12 (3): 199–232.

Gordy, M., and E. Lütkebohmert. 2007. "Granularity Adjustment for Basel II." Deutsche Bundesbank Discussion Paper No. 01/2007.

Glasserman P. 2005. "Measuring marginal risk contributions in credit portfolios". Journal of Computational Finance. Volume 9/Number 2, Winter 2005/06.

Hallerbach W. 2002. "Decomposing portfolio value-at-risk: a general analysis". Journal of Risk. Volume 5, Number 2. Winter 2002/3.

Heyde C, S Kou and X Peng. 2006. "What is a good risk measure: Bridging the gaps between data, coherent risk measures and insurance risk measures", mimeo, Columbia University.

Huang X, H Zhou and H Zhu. 2009. "Assessing the Systemic Risk of a Diversified Portfolio of Asia-Pacific Banks during the Recent Financial Crisis", mimeo.

Hull J. 2006. "VaR vs expected shortfall". Risk. December

Hull J and A White. 2004. "Valuation of a CDO and an n^{th} to Default CDS Without Monte Carlo Simulation," Journal of Derivatives, 12, 2 (Winter 2004), 8-23

International Monetary fund. 2009. Global Financial Stability Report. April

International Monetary fund. 2008. Global Financial Stability Report. April

Jorion P. 2001. Value at Risk: The New Benchmark for Managing Financial Risk. McGraw-Hill.

Knight A D. 2006. "Marrying the micro- and macroprudential dimensions of financial stability: six years on", BIS Speeches. October.

Koyluoglu H U and J Stoker. 2002. "Honour your contribution". Risk. April

Kuritzkes A, T Schuermann, S Weiner. 2007. "Deposit Insurance and Risk Management of the U.S. Banking System: What is the Loss Distribution Faced by the FDIC?". Journal of Financial Services Research. 27:3. 217-242.

Kurth A and D Tasche. 2003. "Contributions to Credit Risk". Risk. March.

Larosière J de. 2009. "The high-level group report on financial supervision in the EU."

Martin R and T Wilde. 2002. "Unsystematic Credit Risk". Risk15. 123-8.

Mas-Colell A, A Whinston and J Green. 1995. Microeconomic Theory. Oxford University Press.

Praschnik J, G Hayt and A Principato. 2001. "Calculating the Contribution". Risk. October

Shapley L. S. 1953. "A Value for n-person Games". In Contributions to the Theory of Games, volume II, by H Kuhn and A Tucker, editors. Annals of Mathematical Studies v 28, pp 307–317. Princeton University Press.

Tarashev N. 2009. "Measuring portfolio credit risk correctly: Why parameter uncertainty matters". BIS Working Paper, # 280.

Tarashev N and H Zhu. 2008. "Specification and calibration errors in measures of portfolio credit risk: The Case of the ASRF model". International Journal of Central Banking. June.

Appendix: Formal results on the impact of size on systemic importance

This appendix provides formal analytical results on the non-linear relationship between the size of an institution and its contribution to system-wide risk, as captured by the Shapley value methodology. More specifically, it proves that, if two institutions are identical in all aspects but size, then the Shapley value of the larger institution divided by that of the smaller one is at least as large as the ratio of the respective sizes.

All results are based on a common framework for the measurement of the risk of a system or subgroup of banks. Risk is driven exclusively by losses related to the failure (default) of individual banks. Given the assumption of a constant loss-given-default (LGD), the loss in the case of a failure of bank i is a constant proportion of the size of the bank: $LGD \cdot S_i$. Then, in addition to the size of each bank, the characteristics that drive its riskiness are: (a) the unconditional probability that it defaults, $PD_i = \text{Prob}\{\text{default } i\}$; and (b) the set of conditional probabilities that i defaults given the default of any group $\{G\}$ of other banks, $PD_{i,G} = \text{Prob}\{\text{default } i \mid \text{default by all } j \in G, i \notin G\}$. The set of conditional PDs would capture any interdependency across banks, stemming from potential "domino effects" (chains of losses across banks) when banks are related via a network of interbank exposures or from the intensity of exposures to common risk factors.

The chosen risk metric is expected shortfall (ES), which equals the expected loss from a given group of banks, conditional on a set of tail events. In turn, a tail event is a loss configuration that delivers extreme aggregate losses. In line with the discussion in Section 2.4, we consider two different types of sets of tail events. A set of the first type is constant for all subgroups of banks and is comprised of tail events in which losses equal or exceed a given quantile of the distribution of losses in the *entire* system. Hence, the expected losses for any subgroup are calculated over the events defined at the level of the entire system. By contrast, a set of the second type is defined at the level of *each subgroup* of banks. In this case, a tail event is defined with respect to the distribution of losses in the subgroup in focus.

In terms of the notation used in the main body of the paper, the fixed set of systemic events gives rise to characteristic function \mathfrak{G}^{II} , while the subgroup-specific set refers to characteristic function \mathfrak{G}^{I} .

Let T be the relevant set of tail events e : $e \in T$. Associated with T there is a set of probabilities $p_{e \in T}$ for the constituent tail events. The ES of a generic group of banks $\{G\}$ can then be expressed as:

$$ES(\{G\}) = \sum_{i \in G} \sum_{e \in T} p_e \cdot LGD \cdot S_i \cdot 1_{[i \in e]}, \text{ where } 1_{[i \in e]} = \begin{cases} 1 & \text{if } i \text{ participates in event } e \\ 0 & \text{otherwise} \end{cases} \quad (\text{A.1})$$

The following two theorems prove results related to the convex positive relationship between a bank's size and its Shapley value. Theorem 1 refers to a constant set of tail events (characteristic function \mathfrak{G}^{II}), while Theorem 2 refers to the case where this set is specific to each subgroup of banks (characteristic function \mathfrak{G}^{I}).

Theorem 1 (characteristic function: \mathfrak{G}^{II}):

Consider two banks S and B , which differ in size, $s_s < s_b$, but have the same risk

characteristics: $PD_s = PD_b$, $PD_{s,b} = PD_{b,s}$ and $PD_{s,j} = PD_{b,j}$. Then $\frac{ShV(B)}{ShV(S)} \geq \frac{s_b}{s_s}$.

Proof of Theorem 1

In the present case, the set of tail events, T , is the same for all subgroups of banks, which simplifies greatly the Shapley value calculation. Given (A1), the marginal contribution of an individual bank i to the risk of a generic subgroup $\{G\}$ equals:

$$ES(\{G, i\}) - ES(\{G\}) = \sum_{e \in T} p_e \cdot LGD \cdot S_i \cdot 1_{[i \in e]}$$

which reflects the fact that T is the same for both subgroups $\{G, i\}$ and $\{G\}$. Note that this marginal contribution is the expected loss associated with i across all tail events (defined at the level of the entire system) and is constant across all subgroups $\{G\}$. This implies that it would also be equal to the Shapley value of bank i , since the latter is a weighted average of such marginal contributions (see Section 2.3 above).

The ratio of the Shapley values of B and S is then given by:

$$\frac{ShV(B)}{ShV(S)} = \frac{\sum_{e \in T} p_e \cdot LGD \cdot S_b \cdot 1_{[b \in e]}}{\sum_{e \in T} p_e \cdot LGD \cdot S_s \cdot 1_{[s \in e]}} = \frac{S_b}{S_s} \cdot \frac{\sum_{e \in T} p_e \cdot 1_{[b \in e]}}{\sum_{e \in T} p_e \cdot 1_{[s \in e]}} \geq \frac{S_b}{S_s} .$$

The reason for the inequality is the following. For each tail event, $e \in T$, that includes S but not B , there must be a corresponding event in T that includes B but not S and has the same probability of occurrence as the former event. This follows from the definition of the set of tail events, T , the size difference, $s_s < s_b$, and the assumption that S and B have identical conditional default probabilities. However, since $s_s < s_b$, it is possible that: (i) there are tail events that include B but not S and (ii) there is no corresponding event that includes S but not B . This implies that $\sum_{e \in T} p_e \cdot 1_{[b \in e]} \geq \sum_{e \in T} p_e \cdot 1_{[s \in e]}$, which establishes the above inequality and completes the proof of the theorem. ■

Theorem 2 (characteristic function: \mathcal{G}^1)

Consider two banks S and B , which differ in size, $s_s < s_b$, but have the same risk characteristics: $PD_s = PD_b$, $PD_{s,b} = PD_{b,s}$ and $PD_{s,j} = PD_{b,j}$. Let S and B have a positive marginal contribution to each subgroup $\{G\}$ of other banks: $ES(\{G, i\}) - ES(\{G\}) > 0$, $i = S$ or B . Then, the following is a sufficient condition for the relative

systemic importance of bank B to be larger than its relative size, ie for $\frac{ShV(B)}{ShV(S)} \geq \frac{s_B}{s_S}$:

$$1) ES(\{i, G\}) - ES(\{G\}) \geq ES(\{i, j, G\}) - ES(\{j, G\}), \text{ where } i, j \in \{S, B\} \text{ and } S, B \notin \{G\}.$$

This condition states that the marginal contribution of bank i to the ES of a subgroup should not decrease as the number of other banks in this subgroup increases. The condition is intuitive because, as the number of banks in the subgroup increases, idiosyncratic risk is diversified away and the impact of each individual bank on the (average) severity of tail events should be expected to decrease. The condition could also be seen as a generalisation of the sub-additivity of ES. Namely, it could be rewritten as $ES(\{i, G\}) + ES(\{j, G\}) \geq ES(\{i, j, G\}) + ES(\{G\})$, which collapses to the sub-additivity property when subgroup $\{G\}$ is empty.

Proof of Theorem 2

The proof incorporates the fact that, under characteristic function \mathcal{G}^1 , tail events differ across subgroup of banks. Concretely, equation (4) above implies that the ratio of Shapley values that is at the centre of Theorem 2 equals:

$$\frac{ShV(B)}{ShV(S)} = \frac{\sum_{G \in \Gamma} \omega(G)(ES(\{B, G\}) - ES(\{G\})) + \sum_{G \in \Gamma} \tilde{\omega}(G)(ES(\{S, B, G\}) - ES(\{S, G\}))}{\sum_{G \in \Gamma} \omega(G)(ES(\{S, G\}) - ES(\{G\})) + \sum_{G \in \Gamma} \tilde{\omega}(G)(ES(\{S, B, G\}) - ES(\{B, G\}))}, \text{ where } \Gamma \text{ is the}$$

set of all subgroups $\{G\}$ that do not contain S or B , and the weights $\omega(G)$ and $\tilde{\omega}(G)$ change

with the number of banks in $\{G\}$. In addition, $\omega(G) \neq \tilde{\omega}(G)$ because, given $\{G\}$, the latter weight is associated with the ES of subgroups comprised of one more bank.

Note next that, given any $\{G\}$ and a marginal contribution $(ES(\{S, B, G\}) - ES(\{S, G\}))$ entering the Shapley value of bank B , there is a corresponding marginal contribution $ES(\{B, G\}) - ES(\{G\})$ that also enters this Shapley value. Similarly for the Shapley value of bank S . This is a result of the Shapley value incorporating the marginal contribution of a bank to *each* subgroup it participates in. Then, the last equality can be rewritten as follows:

$$\begin{aligned} \frac{ShV(B)}{ShV(S)} &= \\ &= \frac{\sum_{G \in \Gamma} (\omega(G) + \tilde{\omega}(G))(ES(\{B, G\}) - ES(\{G\})) + \sum_{G \in \Gamma} \tilde{\omega}(G) \{ [ES(\{S, B, G\}) - ES(\{S, G\})] - [ES(\{B, G\}) - ES(\{G\})] \}}{\sum_{G \in \Gamma} (\omega(G) + \tilde{\omega}(G))(ES(\{S, G\}) - ES(\{G\})) + \sum_{G \in \Gamma} \tilde{\omega}(G) \{ [ES(\{S, B, G\}) - ES(\{B, G\})] - [ES(\{S, G\}) - ES(\{G\})] \}} \\ &\equiv \frac{\Psi^B + \Xi}{\Psi^S + \Xi} \end{aligned}$$

where the fact that the second sum in the numerator is equal to the second sum in the denominator is seen by a simple rearrangement of the summands. Lemma 1, which is stated

and proved below, implies that $\frac{\Psi^B}{\Psi^S} \geq \frac{s^B}{s^S}$. In turn, by condition (1) in the statement of

Theorem 2, $\Xi \leq 0$. Then, since $ShV(S) = \Psi^S + \Xi > 0$ and $s_s < s_B$, it follows that

$$\frac{ShV(B)}{ShV(S)} - \frac{s^B}{s^S} = \frac{\left(\Psi^B - \Psi^S \frac{s^B}{s^S} \right) - \left(\frac{s^B}{s^S} - 1 \right) \Xi}{(\Psi^S + \Xi)} \geq 0. \text{ This proves the theorem. } \blacksquare$$

Lemma 1

Let banks S and B be as specified in Theorem 2 and $\{G\}$ be any subgroup of banks that does

not include either S or B . Then $\frac{ES(\{B, G\}) - ES(\{G\})}{ES(\{S, G\}) - ES(\{G\})} \geq \frac{s_B}{s_s}$.

Proof of Lemma 1

Let $T\{G\}$ denote the set of tail events for a generic subgroup $\{G\}$. Given (A1), for any subgroup of m banks $\{G\}$ we can write $ES(\{G\}) = k \cdot s_G$, where k is a $1 \times m$ vector of probabilities that a bank in $\{G\}$ belongs to the set of tail events, $T\{G\}$, and s_G is the $m \times 1$ vector of respective sizes. Similarly, we can express $ES(\{B, G\}) = ts_B + w \cdot s_G$ and $ES(\{S, G\}) = \hat{t}s_s + \hat{w} \cdot s_G$ where t and \hat{t} are scalars and w and \hat{w} are $1 \times m$ vectors.

The inequality in the statement of the Lemma can be expressed equivalently as a condition on the sign of the following expression:

$$\begin{aligned} \frac{ES(\{B,G\}) - ES(\{G\})}{ES(\{S,G\}) - ES(\{G\})} - \frac{s_B}{s_S} &= \frac{ts_B + w \bullet s_G - k \bullet s_G}{\hat{t}s_B + \hat{w} \bullet s_G - k \bullet s_G} - \frac{s_B}{s_S} = \\ &= \frac{(t - \hat{t})s_B s_S + (s_B - s_S)k \bullet s_G - (s_B \hat{w} \bullet s_G - s_S w \bullet s_G)}{s_S (\hat{t}s_B + \hat{w} \bullet s_G - k \bullet s_G)} \end{aligned} \quad (\text{A2})$$

The Lemma is true if and only if the last expression is (weakly) positive. Given that bank S has positive marginal contributions, the denominator in (A2) is positive. Thus, it remains to prove that the numerator is (weakly) positive. We note the following fact:

Fact 1: $w \bullet s_G \leq \hat{w} \bullet s_G$. In other words, the portion of the ES that is attributed to failures of banks in $\{G\}$ is smaller in the case of subgroup $\{B,G\}$ than in that of $\{S,G\}$.

The reasoning behind this fact follows along the lines of the proof of Theorem 1. Each tail event that is in the set $T\{S,G\}$ and includes bank S (and possibly banks in $\{G\}$) is matched by a corresponding tail event, in $T\{B,G\}$, in which B replaces S. However, the opposite need not be true: there may be some tail events in $T\{B,G\}$ that feature B (and possibly banks in $\{G\}$) but are not matched by tail events in $T\{S,G\}$. If this is the case, then any such tail event, say $[B, \hat{g}]$, enters $T\{B,G\}$ in the place of tail events in $T\{S,G\}$, denoted by $[\tilde{g}]$, which feature only banks from $\{G\}$.

We can establish two properties of tail events $[\tilde{g}]$. First, the probability mass of $[\tilde{g}]$ in $T\{S,G\}$ is equal to the probability mass of the “replacement” tail events $[B, \hat{g}]$ in $T\{B,G\}$. This is by virtue of the fact that the total probability mass of all tail events is constant. Second, the total size of banks in subgroup $\{\hat{g}\}$ that enter a tail event $[B, \hat{g}]$ in the set $T\{B,g\}$ is at most as large as the size of the banks in the tail event $[\tilde{g}]$. To see why, note that, by definition, the aggregate size of banks in a tail event $[\tilde{g}]$ has to be greater than the corresponding size associated with any loss configuration that is not in the set of tail events $T\{S,G\}$. This would be contradicted if the aggregate size of banks in $\{\hat{g}\}$ were larger than the aggregate size of banks in $[\tilde{g}]$ since, then, the aggregate size of banks in $[S, \hat{g}]$, which is not a tail event in $T\{S,G\}$, would be greater than the aggregate size of banks in $[\tilde{g}]$.

The two properties of tail events $[\tilde{g}]$ establish Fact 1.

In turn, Fact 1 points to a lower bound for the numerator of the ratio in (A2):

$$\begin{aligned} &(t - \hat{t})s_B s_S + (s_B - s_S)k \bullet s_G - (s_B \hat{w} \bullet s_G - s_S w \bullet s_G) \\ &\geq (t - \hat{t})s_B s_S + (s_B - s_S)(k \bullet s_G - \hat{w} \bullet s_G) \end{aligned} \quad (\text{A3})$$

The rest of the proof establishes that the right-hand side of inequality (A3) is non-negative.

First note that $t \geq \hat{t}$. In other words, the probability that B participates in the set of tail events $T\{B, G\}$ is at least as high as the probability that S participates in $T\{S, G\}$. The proof of this inequality is identical to a reasoning behind Theorem 1: since banks S and B have identical risk characteristics but $s_s < s_b$, B participates in at least as many tail events as S . This establishes the weak inequality, which implies that the first summand of the numerator in (A3) is positive.

Then note that $k \cdot s_g \geq \hat{w} \cdot s_g$, or that the ES for subgroup $\{G\}$ is at least as large as the portion of the expected losses in $T\{S, G\}$ associated with banks in $\{G\}$. To see why, note that the probability that any loss configuration associated with subgroup $\{G\}$ (be it in the tail or not) is equal to the sum of the probabilities of two loss configurations when the subgroup is $\{S, G\}$: one is identical to the original configuration and one adds bank S . This reflects the fact that the probability of any loss configuration is independent of the banks that are not in this configuration, even if they belong to the subgroup in focus. Then, an argument similar to that underpinning Fact 1 establishes that: (i) for each tail event that is in $T\{S, G\}$ and involves banks in $\{G\}$ (and thus enters the calculation of \hat{w}) there is a corresponding tail event that is in $T\{G\}$ and features the same banks from $\{G\}$ (which enters the calculation of k); and (ii) the opposite need not be true. This establishes the above inequality, which implies that the second summand of the numerator in (A3) is also positive and, thus, completes the proof of the lemma. ■