Discussion of "Measuring Systemic Risk" by V. Acharya, L. H. Pedersen, T. Philippon and M. Richardson

By Mathias Drehmann[1]

In response to the global financial crisis, many policy makers have called for supplementing microprudential regulation focusing on institution-specific risks with a macroprudential approach taking account of system-wide interactions and externalities (e.g. G20, (2009), FSF (2009)). Broadly speaking, the macroprudential approach can be separated along two dimensions (see BIS (2009)). First, there is a time dimension, as shocks get amplified through time by the procyclical nature of the financial system. Second, there is the cross-sectional dimension: given the high degree of interconnectedness, the failure of one institution may have severe ramifications for other participants in the system. A key policy question at the moment is how to operationalise both dimensions. This paper provides a valuable contribution in this area, as it suggests a way to measure and allocate systemic risk, which could be used to address the cross-sectional dimension.

In the first part of the paper, the authors set out a simple theoretical model to highlight how banks should be taxed in the light of externalities, which arise when banks fail. The optimal tax has two components: a microprudential component, which equals the bank's expected loss conditional on default, and a macroprudential component. The latter is labelled by the authors as the bank's systemic expected shortfall (SES), which captures the capital shortfall below a target level in case of a systemic crisis. SES itself can be separated into two parts: The bank's leverage and the bank's marginal expected shortfall (MES), ie the contribution of the bank to losses arising from the systemic crisis.

Based on this insight, the authors estimate SES for over 100 US financial firms prior to the financial crisis and they assess whether it helps to predict the institutions, which had the worst stock-market performance during the financial crisis from July 2007 until December 2008. The authors find that their measure can explain a substantial part of the realized returns. And they show that it outperforms standard measures of risk such as an institution's beta, its volatility or its expected shortfall. They also provide a wealth of information about the time-series behaviour of SES and its determinants at the level of the firm.

In the last par of the paper, the authors discuss how the optimal tax could be implemented in practice and suggest that a private-public partnership similar to terrorism insurance could be optimal. In my comments, I refrain from discussing this issue. And rather than elaborating further on the rich details of the paper, I want to provoke a fruitful discussion by asking some critical questions, first to the way systemic risk is defined and then how the measure is implemented.

The nature of systemic risk

Measurement without definition is impossible. However, in the literature there is no universally agreed definition of systemic risk, even though it is often linked to externalities arising from bank distress (e.g. Rochet and Tirole (1996) or De Bandt and Hartmann (2000)). The authors also relate systemic risk to externalities. They define systemic risk as

[1] The views expressed in this paper are those of the author and not necessarily those of the BIS.

‘widespread failures of financial institutions or the freezing up of capital markets that can substantially reduce the supply of such critical intermediation. Failures of financial institutions thus have an externality on the rest of the economy’ (p. 1). In short, the authors equate systemic risk with a systemic crisis. Whereas the latter is a particular event, risk more generally is characterized by uncertainty: it is forward looking, assessed over a particular horizon, and measured by attaching probabilities to a range of possible outcomes. It may therefore be more precise to define systemic risk as the likelihood and associated costs – including externalities – of a systemic banking crisis. Such a definition would also be more consistent with the model, which actually captures both aspects of systemic risk.

In the empirical part, the authors define a systemic banking crisis more precisely as states where aggregate capital (assumed to equal the aggregate stock market capitalisation of all institutions) falls by more than a specific amount, in this case 60%. Most would agree that such an event would be considered a systemic crisis. But the paper would benefit from a longer discussion of the economic rational behind this choice and why banking capital and stock market capitalisation are the same. But it is also important to assess whether the authors’ definition is sufficiently broad to capture all facets of systemic risk.

To highlight a potential problem, let us assume that two different scenarios hit a banking system of 100 identical banks, each with a tier 1 capital ratio of 10%. In the first scenario, 25 banks loose all their capital. In the second, the capital adequacy of all banks drops to 7.5%. In both scenarios, aggregate capital falls by 25%. It is very likely that systemic externalities will be much higher in the first scenario as 25% of the banking system are in default. In contrast, the banking system should function reasonably well in the second scenario (note that capital adequacy of all banks is well above the Basel minimum of 4%).

Another problem is hidden by the narrow definition of systemic risk. What matters for the allocation of systemic risk in this set-up is a bank’s expected capital shortfall in case of a systemic crisis. The probability of the systemic crisis, on the other hand, affects the macroprudential tax of all institutions in the same fashion. Given it is constant, it is therefore not considered in the empirical analysis. This is a result of the model. But it can only hold if either the probability of a systemic crisis is independent of firms’ capital and investment decisions or if each bank has atomistic size. Clearly, the first condition cannot be satisfied. It seems to me, that the authors implicitly assume the second condition, even though it is unclear whether this is a realistic assumption. Many banking systems are highly concentrated. In Switzerland, for example, two banks dominate the market and the default of a single one would already constitute a systemic crisis. But even the US has several major players and it should be the case that changing the capital requirements of any of these firms will change the overall likelihood of a systemic crisis. It seems crucial to me, that a firms’ contribution to the probability of the systemic crisis materialising is also taken into account when allocating systemic risk, even though it is not apparent how this can be easily integrated in the current empirical framework.

Unsurprisingly, the definition of systemic risk has direct consequences for the measurement and allocation of systemic risk. However, Tarashev et al (2009 and 2010) provide a general methodology which can be applied, independent of the definition of systemic risk. They define the systemic importance of an institution as the share of systemic risk attributed to it by its Shapley value (Shapley, 1953). In this context, the Shapley value of an institution equals its average marginal contributions to systemic risk of all possible sub-groups in the system. Given that the sufficient conditions for the construction of the Shapely value are very weak, it can be used in conjunction with any popular measure of systemic risk and encompasses most allocation procedures that have been studied in the literature.

While the methodology is general, the result, how systemically important different institutions are, still depends on the definition of systemic risk. But even with the same definition of systemic risk at the level of the overall system, Tarashev et al (2009 and 2010) show that different perspectives may lead to different answers. One allocation procedure studied by

Tarashev et al (2009 and 2010) provides a measure of the contribution of each institution to systemic risk. It makes full use of the Shapley value methodology by incorporating information on the risk that the institution generates on its own, as well as on the extra amount of risk generated if it were added to any possible group of other institutions in the system. An alternative perspective is to measure systemic importance by the degree to which each institution participates in a systemic event.

To illustrate the different perspectives, let us assume a system where there is a small, but highly interconnected player in the market, for example an institution which is a counterparty to a large fraction of other market participants. Because of its importance as counterparty, a significant extra amount of risk is generated, if it were added to sub-groups of other institutions in the system. Hence, its measure of systemic importance would be high under the first perspective. On the other hand, systemic importance under the second perspective is determined by the firm's (expected) contribution to system-wide losses, in case of a systemic event. Given its size, losses stemming from this firm will only represent a small fraction of aggregate losses. Hence, under this perspective the firm would only have a relatively low systemic importance. Note that this second perspective is taken by the authors of this paper, where the systemic event is defined as a systemic banking crisis. This may partly explain why the methodology developed in this paper ranks AIG as number 86 (out of 102) in terms of systemic importance (even though AIG was far from small).

Depending on the context, one or the other perspective may be more appropriate. If, as the authors propose, banks are required to buy insurance against systemic risk and the trigger event for the payout of the insurance is a predefined event like a drop in system-wide banking capital, the second perspective provides the actuarially fair insurance premia. Clearly, there is a question whether a single trigger event can be sufficient. Would no externalities arise if aggregate banking capital drops by only 59%? Presumably, this problem could be dealt with by requiring several insurance contracts with different trigger points. However, there is a deeper issue. Let us assume that the small, interconnected firm of the example above is pivotal to the system; if it fails – even for idiosyncratic reasons – many other institutions would fail and the systemic event would materialize. As we have seen, the small player would have to pay very low insurance premia in contrast to large counterparties, which would have to pay significant amounts. It is unclear whether such a scheme would be considered as fair and whether it would provide the right incentives for institutions, like the small firm discussed. It therefore seems that the first perspective may be more appropriate for policy makers, which are concerned about identifying systemically important institutions.

Measuring systemic risk

The implementation of the proposed measure of banks' systemic importance (SES) is relatively simple, as it only requires the estimation of leverage and the marginal expected shortfall (MES) as inputs. Leverage is based on market values. Given that the market value of assets is not observable, the authors use a pragmatic approach and compute quasi market values for assets.[2] The calculation of MES is also straightforward. For data availability reasons MES is initially computed at a 5% confidence level, using daily data on equity returns. Practically, this means the authors take the 5% worst days for the overall market from June 2006 to June 2007, and use each institution's average return on these days as its MES. As mentioned above, the authors assume that aggregate banking capital has to fall by more than 60% before systemic risk chrysalises. Therefore, the authors scale up the MES by

[2] The quasi market value of assets is equal to the book value of assets minus the book value of equity plus the market value of equity.

a factor 60/1.4. They justify this by arguing that during the 5% worst days in the sample, the market fell on average by 1.4% rather than the required 60%.

This scaling method implicitly assumes that all banks participate to the same degree in a loss scenario which is around 40 times more severe than what is measured. This may or may not be the case. In general, extreme loss scenarios for portfolios are driven by exposures which are highly correlated and\or large. It could therefore be the case that this scaling method underestimates the importance of correlations and size. It may actually be cleaner to drop the scaling factor altogether. This would change the weighting of MES and leverage for a bank's systemic expected shortfall. But as it applies to all financial institutions in the same way, it should not affect empirical results materially.

By using daily data, the measurement also assumes that systemic risk crystallizes over a one day period, which the authors discuss only briefly. On the one hand, this assumption may not be overly extreme: the failure of Lehman underlined that banking crises can spread very rapidly. On the other hand, it could be argued that the failure of Lehman was only the culmination of events which started more than a year earlier in August 2007. From this perspective, a one day event window would be too short.

In the broader debate, systemic risk is often associated with the buzzword too big to fail.[3] In this respect it is interesting to note that the measure does not explicitly consider size as an input as SES is a weighted average of MES and leverage. Conceptually, we would expect that larger firms contribute more to the aggregate drop in banking capital. Empirically, the authors show, however, that size and MES are not significantly correlated.[4] The paper only suggests that larger firms can take on more leverage and thereby achieve higher systemic importance. In contrast, other studies find a clear and significant non-linear relationship between size and systemic importance (e.g. see Gauthier et al (2010), Huang et al (2010) or Tarashev et al (2009)).

Notwithstanding these technical issues, the measure could be seen as a first order approximation of the systemic importance of individual financial institutions. Its simplicity may even be a benefit, as the measure can be easily implemented with very little information or computational burden. From a practical perspective, the empirical performance is far more important. A key part of the paper addresses this question, by undertaking an assessment how various risk measures, estimated on data one year prior to the crisis, can explain observed returns during the crisis. The authors show that SES outperforms standard measures of firm-specific risk such as the firm's beta, its volatility or its expected shortfall. It seems, however, that this is primarily driven by leverage. The R^2 is 8.7% if MES is the only explanatory variable. But the R^2 increases to around 24% if leverage is added or the combined measure of systemic expected shortfall is used. This is also apparent from simple correlations presented in the paper, which show that SES is most correlated with event returns, closely followed by leverage (-0.49 versus -0.47). The important role of leverage is further underlined by regression results, based on different one-year estimation periods for leverage and MES starting from January 2006. Leverage is significant for all samples. MES, on the other hand, has only additional explanatory power, if data from May 2007 onwards are included.

---

[3] For example, see The Wall Street Journal, 13 September 2009, "Who's Too Big to Fail? Regulators today won't define 'systemic risk,' unlike 25 years ago"; or "Systemic risk: Are some institutions too big to fail and if so what should we do about it?", committee hearing by the House Committee on Financial Services on 21 July 2009.

[4] This may be may be driven by the definition of "the market". For the purpose of this study, the authors define the market as the value weighted index of all companies covered in CRSP, rather than the value weighted index based on the financial firms covered.

This would suggest that leverage by itself is a good first order approximation of systemic importance. But the result could also be an outcome of a wider phenomenon. Borio and Drehmann (2009a) survey methodologies to measure financial instability. Definitional niceties aside, they show that measures based on market data perform more like "thermometers" of financial stress, in the sense that they provide reliable signals about severity of stress at each point in time. This can be important information for policy makers during a crisis. But to be able to act against the build up of systemic risk, policy makers need "barometers" which can signal pressures with sufficient lead time to implement policy.[5] Most market measures surveyed by Borio and Drehmann increase sharply in August 2007. Some, like the measure of the price of insurance against systemic distress by Huang et al (2009), start signalling pressures only in the first half of 2008, similar to the results found for MES in this paper. But for implementing steps to contain the outbreak of the crisis, this would have been too late.

A general problem with market-based measures is to distinguish between market participants' view of future cash flows and the price they assign to them, ie the risk premium. If the purpose is to identify future distress and systemic vulnerabilities, rather than "prices" attached to them, the influence of the risk premium should be filtered out. This requires several assumptions and is hard to do with any confidence. More importantly, though, any biases in the market's assessment would be embedded in the estimates. If, as some analytical approaches suggest, excessive risk-taking is the source of financial instability, then estimates of systemic risk derived from market prices would tend to be unusually low as vulnerabilities build up ahead of systemic crisis.

Concluding remarks

Addressing all the issues raised in my discussion would be well beyond the scope of this paper or any other single paper. For a start the lack of clear understanding about the nature of systemic risk make measuring and allocating it a formidable challenge. On top of this, practitioners will face empirical difficulties. At this stage, it seems therefore best that policy makers and practitioners should not rely on a single but on a diverse range of tools to measure systemic risk such as simulation models, network approaches, general equilibrium models, simple indicators and the like.[6] The method proposed in this paper could be one of those tools and the conference as a whole a good starting point to explore potential avenues.

---

[5] Borio and Drehmann (2009a) find that leading indicators rooted in the Minsky-Kindleberger view of financial instability appear well suited to identify a build-up of financial instability. Early warning indicators of this type are analysed more closely in Borio and Drehmann (2009b), showing that they would have been fairly successful in providing a signal for several banking systems currently in distress.

[6] Besides the papers in this conference volume, Borio and Drehmann (2009a) provide a recent survey about different approaches to measure financial instability.

Bibliography

Bank for International Settlements (BIS) (2009): *Annual report*.

Borio, C. and M. Drehmann (2009a): "Towards an operational framework for financial stability: "fuzzy" measurement and its consequences", BIS Working Paper 284.

Borio, C. and M. Drehmann (2009b): "Assessing the risk of banking crises - revisited", BIS Quarterly Review, March 29-46.

De Bandt, O. and P. Hartmann (2000): "Systemic Risk: A Survey", ECB Working Paper 35.

Financial Stability Forum (FSF) (2009): Report of the Financial Stability Forum on Addressing Procyclicality in the Financial System.

G20 (2009): "Declaration on strengthening the financial system ", London, 2 April 2009

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

Huang, X., H. Zhou and H. Zhu (2009): "A Framework for Assessing the Systemic Risk of Major Financial Institutions", BIS Working Paper 281.

Huang, X., H. Zhou and H. Zhu (2010): "Assessing the systemic risk of a heterogeneous portfolio of banks during the recent financial crisis", BIS Working Papers 296.

Rochet, J. C. and J. Tirole (1996): "Interbank Lending and Systemic Risk", Journal of Money, Credit and Banking, 28, 733-62.

Shapley, L. (1953): “A Value for n-Person Games”. In Annals of Mathematical Studies, Vol. 28 (Eds, Kuhn, H. and Tucker, A.); Princeton University Press.

Tarashev, N., C. Borio and K. Tsatsaronis (2009): "The systemic importance of financial institutions", BIS Quarterly Review, September.

Tarashev, N., C. Borio and K. Tsatsaronis (2010): "Allocating Systemic Risk to Individual Institutions: Methodology and Policy Applications".