Chen Zhou
c.zhou@dnb.nl

Nikola Tarashev
nikola.tarashev@bis.org

Looking at the tail: price-based measures of systemic importance[1]

We use tools of extreme value theory to extract information about rare events from market prices. We find that such information contributes materially to measures of banks' systemic importance. These measures exhibit strong and intuitive relationships with simple characteristics of banks' balance sheets and income statements.

JEL classification: G20, G28, C14.

The more systemically important a financial institution, the stricter its regulatory requirements should be, all else the same. A Basel III capital surcharge is one manifestation of this macroprudential philosophy (BCBS (2011)). To apply the philosophy, policymakers need to measure systemic importance – that is, an institution's potential contribution to rare but extreme system-wide losses that damage the real economy (Drehmann and Tarashev (2011), Tarashev et al (2010)).

Analysing rare, extreme losses is always challenging because they relate to the tail of the probability distributions of financial shocks, about which data are scarce. In this article, we present an empirical method that addresses this challenge head on and thus provides a potentially useful input to policy discussions.

To illustrate the empirical method, we measure the systemic importance of banks by analysing *explicitly* the tail properties of financial shocks. Two components of this measure are bank size, which we obtain from balance sheet data, and probability of default (PD), for which we rely on commercial estimates. For the other two components – a bank's loss-given-default (LGD) and tendency to default with other banks – we resort to changes in CDS spreads, which reflect shocks to banks' creditworthiness (Jorion and Zhang (2007)). It is tail realisations of such shocks that drive rare but extreme losses in the banking system.

For the analysis of shocks to banks' creditworthiness, we use tools from extreme value theory (EVT). The fundamental idea behind EVT is to focus exclusively on extreme observations in the data, evaluate them and attribute their properties to the *unobserved* tail of shocks' probability distribution. EVT has been applied in

1 The views expressed in this article are those of the authors and do not necessarily reflect those of the Netherlands Bank or the BIS. We thank Angelika Donaubauer for excellent help with the graphs and tables, and Claudio Borio, Stephen Cecchetti, Kostas Tsatsaronis and Christian Upper for useful comments on earlier drafts.

analyses of natural disasters and, since the 1990s, financial crises. The contribution of this article is to employ EVT tools in estimating banks' LGD and tendency to default with others as components of a measure of systemic importance.

By focusing on extreme observations, EVT tools should extract better information about unobserved tail events than alternative approaches that fit *typical* observations to well known probability distributions. But do these tools deliver materially different conclusions? We argue that they do. We find that the EVT-based method and a popular alternative approach disagree substantially on the ranking of banks according to systemic importance.

To further assess the value added of EVT in our context, we examine the extent to which changes in banks' size, PD, LGD and tendency to default with others alter measured systemic importance. We find that the impact of the last two components, which we estimate with EVT tools, is economically significant and similar to that of the first two components, which we obtain directly from the data. In addition, the impact of size, PD and LGD is statistically significant at a high confidence level.

Finally, we examine whether simple balance sheet and income statement characteristics, which are manifestations of banks' business models, can help explain EVT-based measures of systemic importance. We find that the more systemically important banks in our sample tend to be larger, more leveraged and more active in the interbank market than their peers. By contrast, the banks of lesser systemic importance tend to have a higher share of net interest income in total income, resort more to stable sources of funding and exhibit greater operational efficiency.

In the rest of this article, we first define our measure of systemic importance and outline its derivation, paying particular attention to alternative approaches to assessing the tail properties of financial shocks. We then evaluate the extent to which different drivers contribute to the evolution of banks' relative systemic importance over time. Finally, we investigate whether simple bank characteristics can help explain banks' relative systemic importance at particular points in time.

Measuring systemic importance

Before we measure the systemic importance of individual institutions, we need to measure system-wide risk. We equate this risk with the expected credit losses on banks' debt in systemic events, in which losses are large enough to impair financial intermediation and potentially damage the real economy. Concretely, we define a systemic event as one in which the aggregate credit losses exceed a certain fraction of banks' aggregate debt. We abstract from the risks faced by banks' equity holders on the assumption that equity is loss-absorbing and, unless fully depleted, ensures a bank's proper functioning.

Systemic importance is a bank's share in system-wide risk. We define this share to be equal to the expected losses of the bank's creditors in a systemic event. According to this definition, the sum of systemic importance measures across banks is exactly equal to the measure of system-wide risk.

Systemic importance increases with the magnitude of the losses that a defaulting bank gives rise to and with the likelihood that it defaults in a systemic event. Concretely, the larger a bank's debt, the greater the losses it imposes on creditors and the greater the likelihood that its default leads to a systemic event. Henceforth, we refer to the size of a bank's debt as the bank's size. Thus, a bank of a larger size would be of greater systemic importance, all else the same. For similar reasons, the larger the portion of a bank's debt lost in default – ie the larger the *LGD* – the greater the bank's systemic importance. In turn, a bank with a higher (unconditional) probability of default, or *PD*, is more likely to default in a systemic event, all else the same. And this likelihood increases further with the probability that extreme adverse shocks affect the bank at the same time as other banks, or as the *tendency to default with others* increases. Thus, a bank's systemic importance increases with its PD and tendency to default with others.

In the appendix, we derive that the systemic importance of bank i, or SI_i, is the product of three terms:

$$SI_i = size_i \cdot LGD_i \cdot Pr(bank\ i\ defaults \mid systemic\ event) \qquad (1)$$

The last term is the probability of default given a systemic event, or *PDS*. Each of the systemic importance components discussed above – size, PD, LGD and tendency to default with others – affects PDS.

Data and methodology

Our sample consists of 50 large banks headquartered in different parts of the world. Specifically, these are the top 50 banks in terms of total assets (as reported by Bankscope for 2011) for which, in addition to balance sheet data, there are also CDS data (from Markit) and data on expected default frequencies (EDFs, from Moody's KMV). These institutions include 24 European, eight US, five Japanese and four Australian banks, as well as nine banks from emerging market economies.

We think of the 50 banks as forming a system and calculate yearly measures of systemic importance for each one from 2007 to 2011. For these measures, we use the data as follows. First, constrained by data limitations, we set size to be equal to the bank's total non-equity liabilities, net of derivative liabilities, at the end of the year in focus. Ideally, however, the measure of size would have incorporated derivatives positions, as they provide important information about the repercussions of bank defaults.[2] Second, we use one-year EDFs as estimates of banks' unconditional PDs. Third, we combine such PDs with data on CDS spreads, observed over a two-year period ending at the end of the year in focus, to estimate a bank's LGD and its tendency to default with other banks (see below).

Once we have estimates of the four components – size, LGD, PD and tendency to default with others – we proceed as follows. We use the size and LGD estimates as the first two terms of a bank's systemic importance in equation (1). Then, following Huang et al (2009), we define systemic events to be those in which the aggregate default-related losses (eg writedowns) on banks' debt exceed 15% of the

[2] Our methodology would accommodate such a measure of size. In this article, however, we abstract from derivatives as they are reported differently by banks following different accounting rules. Since our analysis focuses exclusively on banks' *relative* systemic importance, the conclusions are affected to the extent that the share of derivatives positions in total non-equity liabilities differs materially across banks.

overall size of this debt in the system.[3] On the basis of this definition, we use banks' sizes and LGDs to identify the systemic events. Next, using the estimates of PDs and tendencies to default with others, we derive the probability of the systemic events in which a particular bank defaults and the probability of all systemic events. The ratio of the first to the second probability is the bank's probability of default in a systemic event, which is the third term in equation (1). Finally, the product of the three terms is our measure of systemic importance.

We next outline the estimation of banks' LGD and tendency to default with others. To fix ideas, we start with the ideal case of plenty of observations from the tail of financial shocks' distribution, which prompts a straightforward estimation procedure. We then turn to the realistic case of no such observations, which calls for using EVT tools.

Ideal case: observations in the tail of interest

Suppose that we observe daily shocks affecting the creditworthiness of each bank in the system. A standard expositional device of the credit risk literature is to identify a bank-specific threshold value for these shocks, beyond which the bank is in default. If we assume that there are plenty of observations of shocks' exceeding the default threshold of each bank in the system, then a PD estimate for a given bank helps to locate the corresponding default threshold. Namely, the location needs to be such that the PD estimate is equal to the share of shocks exceeding the threshold in the total number of shocks affecting the bank.

We think of LGD as reflecting a bank's distress when in default. This makes it natural to estimate LGD as the average distance between adverse shocks that exceed the default threshold and the threshold itself. Since LGD is the share of a bank's debt lost in default, its estimate should be roughly zero if shocks barely exceed the threshold, and should approach 100% for very large shocks. To implement this idea, we would focus on the observed shocks that exceed the threshold. Then we would measure LGD as the average difference between each of these shocks and the threshold, divided by the average value of the same shocks.

To estimate the tendency of a bank to default with a group of other banks, we would need to consider the cases in which all banks concerned experience extreme adverse shocks. Specifically, we would first obtain the number of days for which the shocks affecting *each* bank in the group surpass the corresponding default threshold. Then, focusing on the same days, we would obtain the number of days when the shocks affecting the bank we are interested in also exceed the corresponding threshold. The ratio of the second to the first number of days would be our estimate of the tendency of a bank to default with others.

[3] In practice, policymakers would need to determine the level of system-wide losses beyond which an event would be considered systemic. Given our parameterisation, the default of the eight largest banks in the sample would constitute a systemic event. Admittedly, the most recent financial crisis featured fewer outright defaults by large banks, but this was essentially the result of massive public sector interventions that kept many distressed institutions afloat. Policymakers need to factor out such interventions in assessing banks' systemic importance.

Real-world case: no observations in the tail of interest

Since the shocks affecting banks' creditworthiness are unobservable, we approximate them with daily *changes* in banks' CDS spreads (Jorion and Zhang (2007)). Markets' appetite for risk, which introduces noise in risk estimates, affects such high-frequency changes to CDS spreads to a lesser degree than the spreads' *levels*. Provided that risk appetite evolves slowly over time, we further limit its role by using a short, two-year sample period for CDS spread changes.

However, this sample period is too short to reveal any direct information about shocks that exceed default thresholds, ie about the tail of interest. The highest PD in our sample implies that, if shocks are daily and independent over time, they would surpass the corresponding default threshold with a probability of 0.13%. The two years of daily CDS spread changes we use provide us with roughly 520 observations of such shocks. Since these observations do not reveal direct information about events occurring with probability of less than 0.2% (= 1 / 520), extrapolation is unavoidable. We now compare and contrast two alternative extrapolation methods.

One of the methods is based on extreme value theory (EVT). The fundamental idea underlying EVT is that extreme observations are representative of the tails of the underlying probability distribution (de Haan and Sinha (1999)). We combine EVT-based tools with PD estimates to select *observed* thresholds of CDS spread changes, beyond which the data points are representative of the shocks exceeding the *unobserved* default thresholds for financial shocks. At an intuitive level, the lower the PD, the fewer and more extreme the selected data points.[4] In Graph 1 (left-hand panel), where we plot the daily CDS spread changes for two banks over two years, EVT tools indicate that the observations in red carry information about the tail of interest.

Once we have identified the observed thresholds and the data points that exceed them, we implement the estimation algorithm outlined in the previous subsection. Thus, large differences between the selected data points and the corresponding threshold would lead to a high LGD estimate, as they imply a high likelihood that the losses on the bank's debt would be severe.[5] Likewise, selected data points that have a strong positive correlation across banks would result in high estimates of banks' tendency to default with others.

The alternative approach we consider assumes that the shocks driving banks into default are normal random variables (Li (2000)). It calls for estimating the sample mean and variance-covariance matrix of the CDS spread changes. These two estimates, which reflect mainly information about typical shocks, pin down a normal distribution. The approach uses the analytic formulas characterising this distribution to extrapolate as far into the tail as desired.

To compare the implications of the two approaches, we simulate two years of data from a normal distribution that has the same mean and variance-covariance matrix as the actual CDS spread changes for the two banks in Graph 1. We plot the

4 Choosing more data points reduces the random noise in our tail estimate but also brings in more information about the centre of the distribution, thus introducing a bias in the tail estimate. To obtain an optimal mix between random estimation noise and bias, we rely on the selection methodology developed in de Haan and de Ronde (1998).

5 We rescale our LGD estimates so that they average 50% across banks and years. This is a frequently used value in the credit risk literature. The higher the average LGD, the lower the number of defaulting banks that would generate a systemic event, all else the same.

Two approaches to estimating the tail

In basis points | Graph 1

EVT approach[1] | Normal-based approach[2]

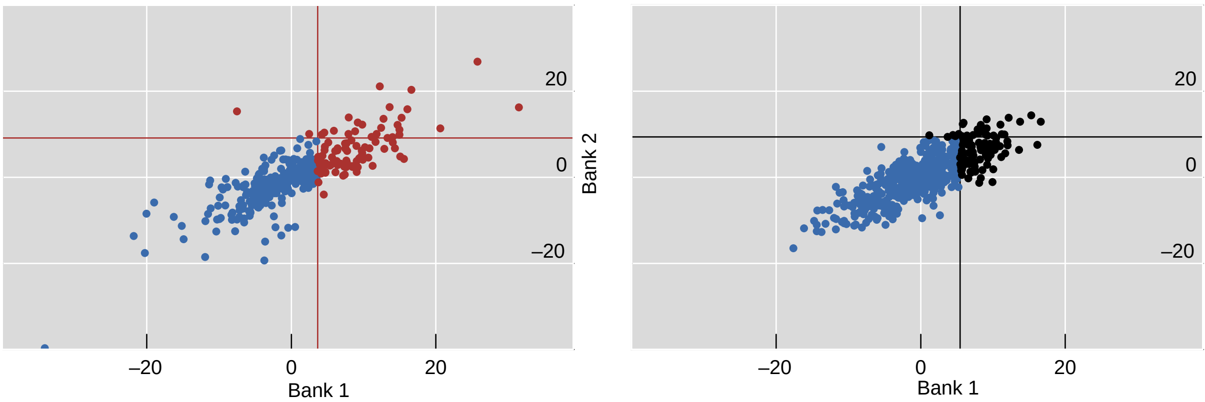

The red and black dots correspond to the same percentiles of the respective distributions. The vertical (horizontal) reference lines correspond to the observed thresholds for bank 1 (bank 2).

[1] Daily changes in CDS spreads between 1 January 2010 and 31 December 2011, for two banks in the sample. The red dots indicate the data used by the EVT approach. [2] Simulated data, drawn from a normal distribution that has the same mean and variance-covariance matrix as the data in the left-hand panel.

Sources: Markit; authors' calculations.

outcome in the right-hand panel and flag extreme observations with black dots. A casual comparison between the two panels reveals that the black dots associated with the normal distribution are clustered more closely to the default thresholds and are less synchronised across the two banks than the red dots associated with the actual data. In the light of the above discussion, it then comes as no surprise that the normal distribution leads to an estimated probability of the two banks defaulting together that is less than two thirds of the corresponding estimate under the EVT approach. This is a general phenomenon. The ratios between the normal-based and EVT estimates of the probability of joint defaults average about one third across all pairs of banks in our sample and are smaller than unity for 95% of the pairs.

In contrast to approaches that rely on the analytical formulas describing *entire* distributions, the EVT approach focuses exclusively on the tails of interest. Of course, it could happen that typical shocks describe these tails well and, thus, the EVT and normal-based approaches lead to similar conclusions. In other cases, the EVT approach could agree with an approach based on a distribution that implies a high probability of extreme shocks and high degree of synchronisation of such shocks across institutions, such as a Student's *t*-distribution. What sets EVT apart is that it is flexible enough to seamlessly account for tail properties that differ across banks and change over time.

Empirical results on systemic importance

Our analysis of the measures of systemic importance has four takeaways. First, the conclusions reached by applying the EVT approach differ materially from those obtained on the assumption that shocks have a normal distribution. Second, even though we can confidently differentiate many banks according to their EVT-based measures of systemic importance, substantial uncertainty remains. Third, each component of these measures – size, PD, LGD and the tendency to default with others –materially affects the evolution of banks' relative systemic importance over time. Fourth, simple bank characteristics help to explain the variation of EVT-based measures of systemic importance in the cross section.

We focus the analysis on the way in which the price-based measures of systemic importance differentiate banks, either through their ranking or their *relative* systemic importance. Our choice reflects existing evidence that markets have done a reasonable job in differentiating banks ex ante with respect to their losses in a financial crisis (Acharya et al (2009)). This contrasts with markets' failure to accurately assess the actual level of system-wide risk and, thus, the *level* of individual institutions' systemic importance (BIS (2011), Chapter VI).

EVT versus normal-based ranking of banks

Do different empirical approaches to evaluating the tail properties of the data rank banks differently in terms of their systemic importance? To answer this question, we compare banks' rankings under the normal-based approach to those under the EVT approach. In a given year, we obtain the absolute difference between the two alternative rankings for each bank and then take an average in the cross section. The resulting average ranking change (ARC) is substantial. For the entire sample of 50 banks over five years, ARC ranges from roughly five positions in 2011 to seven positions in 2008. Similarly, for the 25 most systemically important banks according to the EVT approach, ARC varies between roughly five positions in 2007 and seven positions in 2009.

Given that the EVT- and normal-based approaches estimate an *unobservable* tail of the shocks' probability distribution, it is not possible to test which approach delivers more accurate results. A priori, however, the extreme observations on which the EVT approach focuses provide better information about the tail properties of financial shocks than the typical observations that strongly affect the conclusions of the normal-based approach. Thus, we analyse only EVT-based measures in the rest of this article.

Uncertainty around banks' rankings

Can we confidently rank banks with respect to measured systemic importance? To address this question, we quantify the estimation noise around each point estimate and the correlation of this noise across banks (see Hartmann et al (2004)). Then, considering one bank at a time, we determine how many other banks are of statistically higher or lower systemic importance. Graph 2 reports the results for 2011, ordering banks from high to low systemic importance.

Graph 2 provides four pieces of information about each bank. First, the black dots indicate each bank's ranking according to the point estimate of its systemic importance. Second, the height of the white bar corresponds to the number of other banks whose systemic importance is statistically indistinguishable, with 95% confidence, from that of the bank in focus. Third, the height of the blue bar denotes the number of other banks that have significantly lower systemic importance. Fourth, the height of the red bar indicates the number of other banks with significantly higher systemic importance. In the absence of estimation noise, the blue bars would reach up to the corresponding black dots, the red bars would extend down to these dots and the white bars would disappear.

Reassuringly, we can differentiate many banks with a high level of confidence. For instance, the systemic importance of all but one of the banks in the sample is statistically different from that of 28 or more other banks. That is, virtually every bank can be differentiated from more than half of the others in the sample.

Admittedly, however, the differentiation across banks is far from perfect. For example, the banks ranked 13 and 15 can be confidently distinguished from only one other bank with a higher point estimate of systemic importance (see the corresponding red bars). In other words, the systemic importance of banks 13 and 15 is potentially much higher than what point estimates indicate. Conversely, bank 11 may be ranked too high, as 21 other banks could in fact be of greater systemic importance (see the combined height of the corresponding red and white bars). Thus, when interpreting results on systemic importance, careful attention must be paid to estimation noise.

Finally, putting estimation noise in the background, we find that banks' relative systemic importance varies considerably over time. To illustrate this, we calculate the systemic importance as a share in system-wide risk for the top 10 banks in 2011 (Graph 3). The most stable of these shares, that of bank 3, varies between 0.033 in 2007 and 0.053 in 2011, a range that amounts to almost 50% of the average share of this bank over the sample period. At the other extreme, bank 9 saw its share in system-wide risk range between 0.015 in 2008 and 0.084 in 2010, which amounts to 150% of the bank's average share over the sample period. This variability prompts us to examine the relative strength of the four components of our systemic importance measure – the two we obtain directly from the data and the two we estimate with EVT tools – as drivers of the measure's evolution over time. We do this in the next subsection.

Differentiating banks according to systemic importance

With 95% confidence[1]

Graph 2

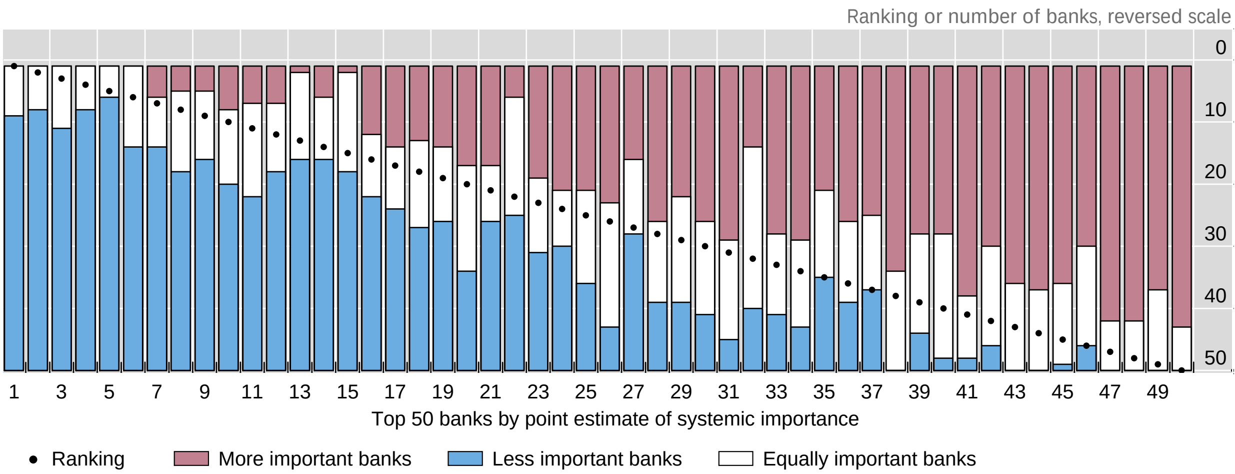

[1] Based on bootstrapped confidence intervals around point estimates for 2011. The height of the red (blue) bars corresponds to the number of significantly more (less) important banks. The height of the white bars corresponds to the number of other banks whose systemic importance cannot be distinguished from that of the particular bank at the chosen confidence level.

Sources: Bankscope; Markit; Moody's KMV; authors' calculations.

Systemic importance over time[1]

Shares in system-wide risk, in per cent — Graph 3

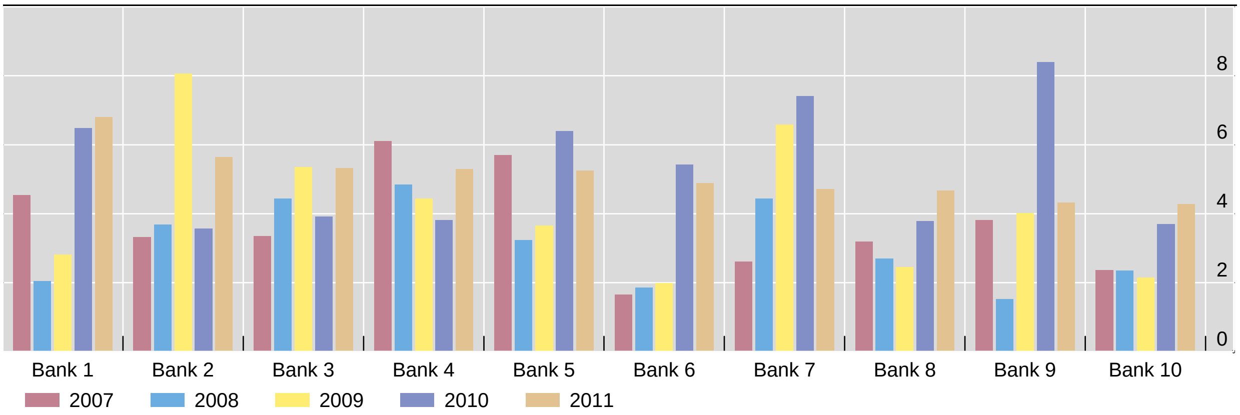

[1] Banks' numbers correspond to their 2011 rankings in terms of systemic importance.

Source: Authors' calculations.

Drivers of banks' relative systemic importance

We examine the extent to which the four components – a bank's size, PD, LGD and tendency to default with others – drive changes in banks' relative systemic importance from 2009 to 2011.[6] To this end, we derive what the estimates of banks' systemic importance would have been had one of the drivers remained as in 2009 while the other three changed to their 2011 levels. In Graph 4, we use diamonds to plot these hypothetical estimates for five banks, holding a different driver fixed in each panel. For instance, the diamond for bank 1 in the left-hand panel shows what this bank's systemic importance would have been in 2011 had it kept its 2009 size. In addition, we plot the corresponding *actual* estimates of systemic importance for 2009 (red dots) and 2011 (blue dots), which do not change across panels. In each case, we express systemic importance as a share in system-wide risk.

We find that each driver has a strong impact on the evolution of the point estimates of banks' relative systemic importance. This result reflects a comparison between two sets of rankings: one based on the actual measures for 2011, and one on the hypothetical measures. For each driver, we calculate the average ranking change (ARC) between the two sets of rankings and report the results in Table 1 (left-hand columns). Keeping any of the drivers as in 2009 leads to an ARC of seven to eight positions (first column). And this effect remains strong for the 25 most systemically important banks in 2011, for which the ARC is six to seven positions (second column).

The general message is similar when we examine the statistical significance of the drivers' impact. For this exercise, we analyse the noise around the actual point estimates of relative systemic importance in 2009 and 2011 and the hypothetical point estimates. We say that a driver has a strong impact on a bank if the actual and hypothetical estimates for 2011 (the blue dots and diamonds in Graph 4) are

[6] The 2009 and 2011 estimates of systemic importance are based on *non-overlapping* two-year samples of CDS data.

Strength of different drivers

In per cent | Graph 4

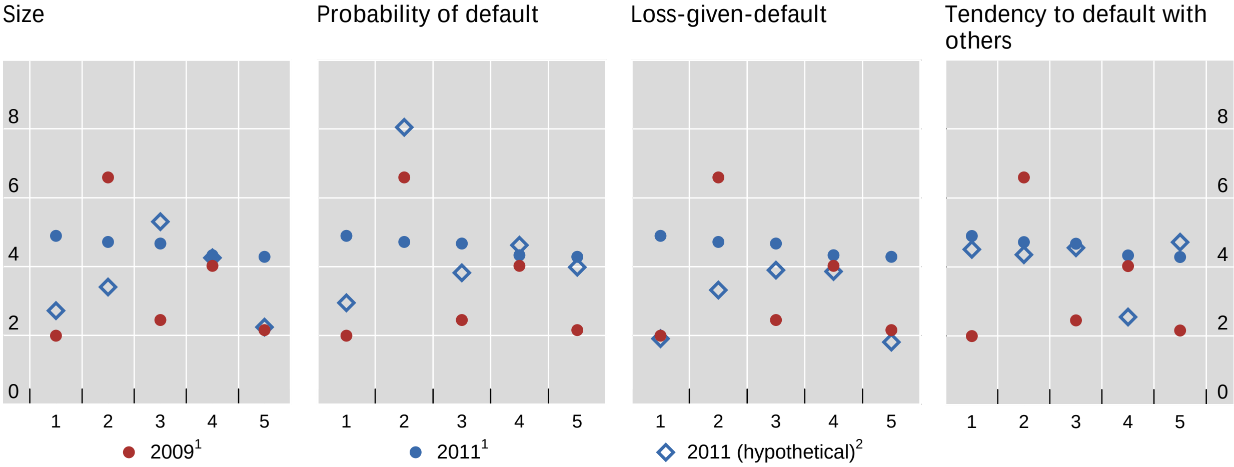

[1] Systemic importance, as a share in system-wide risk, for five banks in the sample. [2] Systemic importance, as a share in system-wide risk, when the driver indicated in the panel heading is held as in 2009 and the other three drivers are as in 2011.

Source: Authors' calculations.

statistically different at the 95% confidence level. As reported in Table 1 (third column), three of the four drivers – size, PD and LGD – have a statistically significant impact for at least half of the banks. And for 21 banks (= 7 + 12 + 2), size, PD or LGD is the only driver to have a statistically significant impact on measured systemic importance (fourth column).[7]

The analysis in this subsection underscores the strength of the drivers we obtained with EVT tools, ie banks' LGD and tendency to default with others. Each of these drivers has a material impact on the point estimates of banks' ranking according to systemic importance. And while the tendency to default with others comes with much estimation noise, LGD has a statistically significant impact on the relative systemic importance of half of the banks in the sample.

Simple bank characteristics and systemic importance

Is there a relationship between our price-based measures of systemic importance and bank characteristics derived from balance sheets and income statements? If a relationship does exist, then it would shed light on which aspects of banks' business

[7] For a driver to qualify as the only one with a significant impact on a particular bank, three criteria must be satisfied. First, the actual estimate of the bank's relative systemic importance in 2011 has to be significantly different from the corresponding hypothetical estimate for which this driver is held as in 2009; ie the driver should have a statistically significant impact. Second, the same hypothetical estimate has to be statistically indistinguishable from the actual 2009 estimate for the bank. Third, the individual impact of each of the other three drivers on the bank should not be statistically significant.

Drivers of systemic importance

Table 1

	Average rank change[1]		Impact on relative systemic importance[2]	
	All 50 banks	Top 25 banks	Significant impact[3]	Only significant driver[4]
Size	8	6	29	7
Probability of default	7	7	39	12
Loss-given-default	8	6	25	2
Tendency to default with others	8	7	1	0

[1] Average of the absolute differences between banks' rankings in 2011 and the corresponding rankings when the driver in the row heading is kept as in 2009 and all other drivers are as in 2011; rounded to a whole number. [2] Number of affected banks, at the 95% confidence level. Based on bootstrapped confidence intervals, expressing systemic importance as a share in system-wide risk. [3] The following condition must be satisfied: keeping the driver indicated in the row heading as in 2009 but letting all other drivers change to their 2011 levels leads to a measure of relative systemic importance that is statistically different from the actual measure for 2011. [4] Three conditions must be satisfied: (i) the driver indicated in the row heading has a significant impact on measured systemic importance; (ii) keeping this driver as in 2009 but letting all other drivers change to their in 2011 levels leads to a measure of relative systemic importance that is not statistically different from the actual measure for 2009; and (iii) the individual impact of each of the other drivers on measured systemic importance is not statistically significant.

Source: Authors' calculations.

models shape markets' perception of systemic importance.[8] We now look for such a relationship in the cross section of banks at each year in the sample.

To pursue our analysis, we focus on banks' probability of default in a systemic event (PDS), the third term of systemic importance in equation (1). There are two reasons for this choice. First, we find that LGD has a negligible impact through the second term on the variation of systemic importance in the cross section of banks. This is despite the key role of LGD as a driver of the evolution of banks' relative systemic importance over time (as described in the previous subsection). Second, the first term – bank size – is directly related to a balance sheet feature by construction, ie non-equity liability net of derivative liabilities.

We investigate the relationship between PDS and six bank characteristics. The first is bank *size*, as defined above, which we consider because a systemic event is more likely to feature defaults by larger banks. The second characteristic is *leverage*, or assets divided by equity, which is a proxy for a bank's credit riskiness and, thus, for its unconditional PD. The third is the *stable funding ratio* – defined as customer deposits and long-term debt over total liabilities – which reflects the extent of funding liquidity risk to which a bank is exposed. The fourth characteristic is the ratio of interbank to total assets and captures *interbank links* that affect the tendency to default with others. The fifth variable is *net interest income* as a share in total net income. Interest income tends to be the most robust income type, thus contributing to banks' resilience at a time of general distress (BIS (2012), Chapter VI). Finally, the *cost-to-income ratio* is inversely related to a bank's efficiency and, by extension, to the capacity to cut costs in order to stay afloat at a time of widespread difficulties. We would expect size, leverage, interbank links and cost-to-income to be positively related to PDS, and the stable funding ratio and net interest income to be negatively related to PDS.

[8] See Ayadi et al (2012) for a broad analysis of business models in banking. Building on that paper, Blundell-Wignall and Roulet (2012) use business model indicators to explain banks' distance-to-default, a measure akin to an unconditional PD.

Simple bank characteristics and probability of default in a systemic event[1] Table 2

	Bivariate relationships[2]					Multivariate regression: decomposing the goodness of fit[3]				
	2007	2008	2009	2010	2011	2007	2008	2009	2010	2011
Size	0.29**	0.19	0.25*	0.47***	0.36**	0.03	0.03	0.05	0.18***	0.08**
Leverage	0.41***	0.07	0.18	0.33**	0.41***	0.09	0.00	0.03	0.07*	0.13***
Cost-to-income	0.41***	0.25*	0.33**	0.43***	0.50***	0.12**	0.06*	0.09**	0.14***	0.21***
Interest income	–0.28*	–0.06	–0.09	–0.24*	–0.26*					
Stable funding	–0.01	–0.21	–0.18	–0.20	–0.33**					
Interbank links	0.18	0.36**	0.13	0.31**	0.30**					
Total R-squared						*0.24*	*0.09*	*0.17*	*0.39*	*0.42*

***, ** and * indicate significance at the 99%, 95% and 90% confidence levels, respectively.

[1] Size = total non-equity liabilities minus derivative liabilities; leverage = total assets minus derivative assets divided by total equity; cost-to-income = operating expenses divided by total net income; interest income = net interest income, as a share in total net income; stable funding = customer deposits plus long-term debt as a share in total liabilities; interbank links = interbank assets as a share in total assets minus derivative assets. [2] Cross-sectional correlation between the variable in the row heading and the probability of default in a systemic event. [3] Obtained from a linear regression of the probability of default in a systemic event on size, leverage and cost-to-income. The top three numbers in each column add up exactly to the fourth.

Source: Authors' calculations.

Table 2 (left-hand columns) reveals that *all* bivariate relationships are of the expected sign, in each year of the sample. In terms of statistical significance, these relationships weaken during the crisis years, 2008 and 2009, but are in general quite strong both before and after. Interestingly, the cost-to-income ratio is the characteristic that most consistently helps to differentiate banks with respect to PDS. By extension, this finding suggests that, when evaluating the likelihood that a bank will fail in a systemic event, markets appear to pay closer attention to efficiency than to size or leverage. Less surprisingly, stronger interbank links and less reliance on interest income tend to be associated with a high PDS.

We also explore the extent to which several bank characteristics can *simultaneously* explain PDS. Unfortunately, because different characteristics tend to go hand in hand, using them simultaneously in regression analysis makes it hard to distinguish their separate relevance. This leads us to a parsimonious specification in which we explain PDS on the basis of bank size, leverage and cost-to-income ratio, in each of the five years in the sample.

The resulting goodness-of-fit measures (Table 2, right-hand panel) confirm the general message from the bivariate analysis: simple bank characteristics can explain PDS quite well before and especially after the peak of the crisis, but fare quite poorly in 2008 and 2009. The characteristic with the strongest and most robust explanatory power is again the cost-to-income ratio.

Conclusions

Measuring systemic importance involves analysis of rare, tail events, about which relevant data are scarce. In trying to address this issue, we employ tools of extreme value theory to infer the tail properties of financial shocks from market prices. We find that exploiting these properties enhances our understanding of systemic importance and delivers measures that exhibit strong and intuitive relationships with simple bank characteristics.

Of course, the measures of systemic importance we derive paint only part of the picture. They reflect prices from only one market and are limited to publicly available data on banks' balance sheets and credit risk. Thus, they need to be complemented by information from other financial markets and supervisory assessments of banks' riskiness and interconnectedness.

References

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

Ayadi, R, E Arbak and W P de Groen (2012): "Regulation of European banks and business models: towards a new paradigm?", Centre for European Policy Studies.

Bank for International Settlements (2011): *81st Annual Report*.

——— (2012): *82nd Annual Report*.

Basel Committee on Banking Supervision (2011): *Global systemically important banks: assessment methodology and the additional loss absorbency requirement*, consultative document, July.

Blundell-Wignall, A and C Roulet (2012): "Business models of banks, leverage and the distance-to-default", *OECD Journal: Financial Market Trends*, vol 2012, issue 2.

de Haan, L and J de Ronde (1998): "Sea and wind: multivariate extremes at work", *Extremes,* vol 1, no 1 , pp 7–45.

de Haan, L and AK Sinha (1999): "Estimating the probability of a rare event", *The Annals of Statistics*, vol 27, no 2, pp 732–59.

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

Hartmann, P, S Straetmans and C de Vries (2004): "Asset market linkages in crisis periods", *Review of Economics and Statistics*, vol 86, no 1, pp 313–26.

Huang, X, H Zhou and H Zhu (2009): "A framework for assessing the systemic risk of major financial institutions", *Journal of Banking and Finance*, vol 33, no 11, pp 2036–49.

Li, D (2000): "On default correlation: a copula function approach", *Journal of Fixed Income,* vol 9, no 4, pp 43–54.

Jorion, P and G Zhang (2007): "Good and bad credit contagion: evidence from credit default swaps", *Journal of Financial Economics,* vol 84, no 3, pp 860–83.

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

Annex: Defining systemic importance

In this annex, we present a formal definition of our measure of systemic importance. In our analysis we define a systemic event as an event in which the aggregate losses on the debt of all banks in the system exceed some fraction, α, of the size of this debt. Concretely, a systemic event occurs when $\sum_{j=1}^{N} L_j > \alpha \sum_{j=1}^{N} Size_j$. In this expression, an index j refers to a particular bank, L_j is the loss on this bank's debt, $Size_j$ denotes the size of this debt, and N is the total number of banks in the system.

We equate the systemic importance of bank i, or SI_i, with the expected loss on this bank's debt in systemic events: $SI_i = E\left(L_i \middle| \sum_{j=1}^{N} L_j > \alpha \sum_{j=1}^{N} Size_j\right)$. Then, we write the loss L_i as the product of the bank's debt, the fraction of this debt that is lost at default, and an indicator that is equal to one if the bank is in default and zero otherwise: $L_i = Size_i \cdot LGD_i \cdot I_i$. Finally, treating size and loss-given-default (LGD) as parameters, we obtain an explicit version of equation (1) in the main text:

$$SI_i = Size_i \cdot LGD_i \cdot Pr\left(I_i = 1 \mid \sum_{j=1}^{N} L_j > \alpha \sum_{j=1}^{N} Size_j\right).$$