C Systemic risk, contagion and financial networks[168]

This special feature proposes a methodology to measure systemic risk as the percentage of banks defaulting simultaneously over a given time horizon for a given confidence level. The framework presented here is applied to euro area banks. It is observed that since the announcement of the comprehensive assessment in October 2013 banks have significantly reshuffled their security portfolios. This has resulted in a decline in the probability of systemic events occurring.

Introduction

Although widely referred to, the concept of systemic risk remains elusive and hard to quantify (see, for instance, Hansen (2013)).[169] In an attempt to fill this gap, this special feature defines systemic risk as the risk that a "large" number of banks default simultaneously with negative reverberating effects on the real economy. In line with this definition, this article measures systemic risk as the systemic Value-at-Risk of a banking system, i.e. as the percentage of banks going bust simultaneously over a given time horizon for a given confidence level. By reverse engineering, this framework also allows us to evaluate the probability that a systemic event occurs: after setting a percentage of banks failing simultaneously, the probability associated with this event is estimated.

The estimates of the systemic Value-at-Risk and of the probability of a systemic event are derived from a distribution of the yearly number of bank defaults. In the proposed framework, contagion is the factor generating fat tails in this distribution, which allows us to capture systemic risk. In particular, the model characterises contagion through fire sales: if a bank defaults because of an idiosyncratic shock, this failure can contaminate other banks via their common exposures. Failing banks liquidate their security portfolios, transmitting shocks from one bank to another through fire sales.

The distribution of the number of bank defaults is generated with Monte Carlo simulation techniques using data relative to the network of banks' common exposures. This enables the model to capture how the topology of a banking system network affects systemic risk: it is the architecture of such a network that determines how contagion propagates and how resilient the system is. Specifically, this distribution is derived by letting banks fail in line with their idiosyncratic shocks, which can trigger as a consequence a fire sale.

Empirical evidence supports the intuition that systemic risk materialises in parallel with a "large" number of banks failing at the same time and permits us to qualify the notion of "large". When looking at the number of bank defaults in the United States from 1934 to 2014, three episodes stand out: the Great Depression in the 1930s, the

---

[168] Prepared by Lorenzo Cappiello, Linda Fache Rousová and Mattia Montagna.

[169] Hansen, L.P. (2013), "Challenges in Identifying and Measuring Systemic Risk", in Brunnermeier, M.K. and Krishnamurthy, A. (eds.), *Risk Topography: Systemic Risk and Macro Modelling*, NBER Books Series.

savings and loan crisis in the 1980s and the 1990s, and the recent financial crisis in the second half of the 2000s (see the left-hand panel of Chart C.1). It can be safely argued that systemic risk materialised during these three episodes. Next, the yearly percentage of US failing banks is computed. This is done by dividing the number of bank defaults at the end of each year by the number of active banks at the beginning of the year. The right-hand panel of Chart C.1 reports the empirical distribution of this percentage of bank failures. Two key points are revealed by this distribution. First, since systemic events are defined as those episodes characterised by the simultaneous failure of a "large" number of banks, the distribution enables us to qualify the notion of "large". The right-hand panel of Chart C.1 shows that when more than 3% of the total number of banks fail at the same time, a deep financial crisis ensues. Second, the shape of the distribution is fat tailed, which can only be explained by introducing non-zero correlations between banks' default probabilities. Moreover, since the standard deviation of the distribution is equal to 0.8 but the probability mass in its tail is quite large, a Gaussian function cannot describe it. A Gaussian distribution of the number of bank defaults would imply that crises wiping out more than 3% of the US banking system would occur once every 700 years.

Chart C.1

Historical distribution of bank defaults in the United States

(x-axis: number of defaulting banks; percentages)
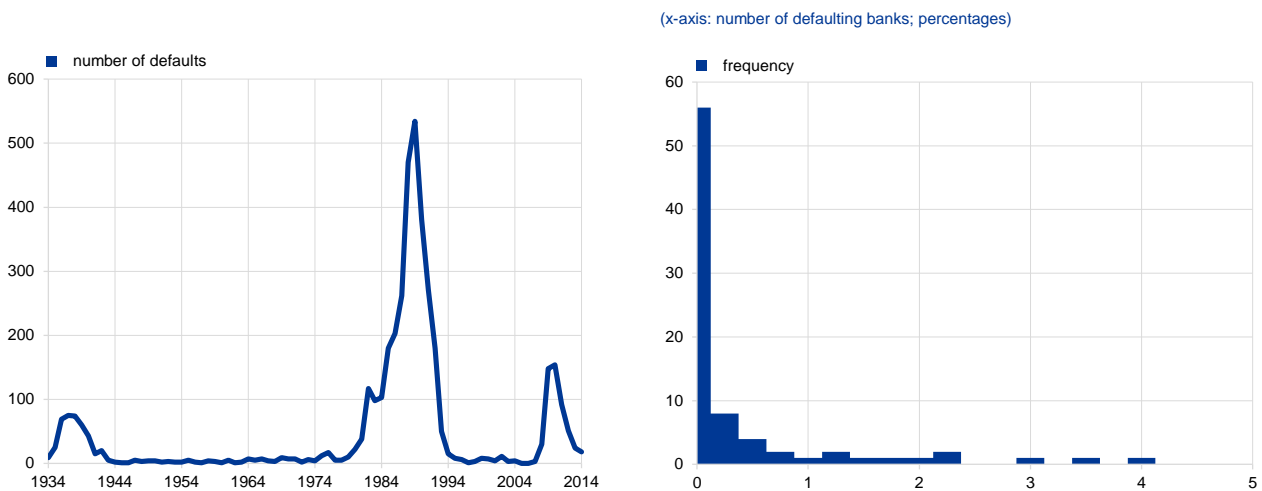
Sources: Federal Deposit Insurance Corporation and ECB calculations.
Notes: The left-hand panel reports the total number of defaults in the United States from 1934 to 2014. The right-hand panel reports the distribution of the yearly percentage of US failing banks computed as a fraction of the total number of active banks.

The approach proposed in this special feature has two main advantages over other methodologies. First, systemic risk is measured without relying on historical data. Instead, its estimate is based on the actual architecture of a banking network and on a simple contagion mechanism. Since systemic events are rare, historical data typically do not contain enough information to make proper inference. Similarly, measures of systemic risk based on past asset prices suffer from the drawback that price developments are cyclical. Asset prices do not necessarily convey information about vulnerabilities well in advance of a crisis, often collapsing just before a crisis

materialises.[170] Second, this framework enables us to isolate the role of contagion in generating systemic risk. Shedding light on the root factors of systemic risk has vast policy implications since it allows policy-makers to adopt the most efficient risk-mitigation measures. Finally, providing a measure to quantify systemic risk can contribute to increasing the accountability of the policies aimed at counteracting it.

Although the concepts discussed in this special feature are sufficiently general and can be applied to the whole financial system, this article focuses on banks. Similarly, although contagion materialises through fire sales for the purpose of the discussion, the framework is general enough to accommodate a variety of contagion models.

Methodology

When deriving the distribution of the number of bank defaults, it is necessary to take into account that banks' default probabilities are correlated. Such correlations produce fat tails in the distribution and therefore constitute the key ingredient to capture systemic events. In this framework, contagion is the main factor generating non-zero correlations between the probabilities of banks' defaults (for a formal sketch of the methodology, see the box).

The specific contagion mechanism which is used in this special feature is a fire-sale model in the spirit of Greenwood et al. (2015) [171] and Eisenbach et al. (2015).[172] In such a model, after an idiosyncratic shock, banks sell off assets to restore their desired target leverage. But sales can depress asset prices, ultimately eroding other banks' capital, which may trigger another bout of sales and may further contract prices and reduce capital, until a new equilibrium is achieved.[173]

The distribution of the number of bank defaults is obtained by adopting the following simulation strategy. At each time *t*, banks are allowed to fail according to their idiosyncratic default probability. Then, it is assumed that each failing bank liquidates its security portfolio, depressing securities' prices and triggering fire sales, which can eventually produce further defaults. But there is more to it than that. In the framework proposed here, the propagation and amplification of shocks, and ultimately the stability of the banking system, will depend on the topology of the network of banks' overlapping portfolios. While shocks to individual banks are transmitted to the whole banking system via the contagion mechanism of the fire sales, such transmission varies according to the architecture of the network of banks' common asset exposures. Thanks to this intuition, the construction of the distribution of the number of bank defaults takes into account the topology of the interbank network.

---

[170] For example, in the week before its demise, Lehman Brothers' senior bonds were rated A by Standard & Poor's and A2 by Fitch. See Giglio, S. (2014), "Credit Default Swap Spreads and Systemic Financial Risk", Working Paper.

[171] Greenwood, R., Landier, A. and Thesmar, D. (2015), "Vulnerable Banks", *Journal of Financial Economics*, Vol. 115, Issue 3, pp. 471-485, March.

[172] Eisenbach, T. and Duarte, F. (2015), "Fire-Sale Spillovers and Systemic Risk", Federal Reserve Bank of New York Staff Report, No 645, February.

[173] See also Cappiello, L. and Supera, D. (2014), *Fire-sale externalities in the euro area banking sector*, *Financial Stability Review*, European Central Bank, November.

The idiosyncratic shocks to individual institutions are by definition uncorrelated. By analysing the effect of such shocks, the framework enables us to isolate the contribution of contagion to systemic risk, since the impact which may be due to other shocks hitting simultaneously all banks' balance sheets are removed.

After constructing the distribution of the number of bank defaults it is possible to estimate the systemic Value-at-Risk and the probability of a systemic event. For instance, setting the probability that a systemic event occurs to a pre-specified probability β (e.g. 5%), the associated Value-at-Risk ($\text{SysVaR}_t(\beta)$) indicates the percentage of banks (e.g. 20%) going bust in such an event (see the left-hand panel of Chart C.2 for illustration). By reverse engineering, if a systemic event is considered to occur where at least s% of banks (e.g. 20% of banks) go bust simultaneously, the framework allows us to estimate the probability $P_m(s)$ that such a systemic event occurs (see the right-hand panel of Chart C.2).

Once the systemic Value-at-Risk is estimated, it is necessary to evaluate whether it is indicative of a systemic event. In line with empirical evidence collected for the US banking system (according to which a simultaneous failure of more than 3% of banks is associated with a deep financial crisis – see Chart C.1), each time the systemic Value-at-Risk is larger than 3%, a systemic crisis can materialise with a probability larger than β.

Chart C.2
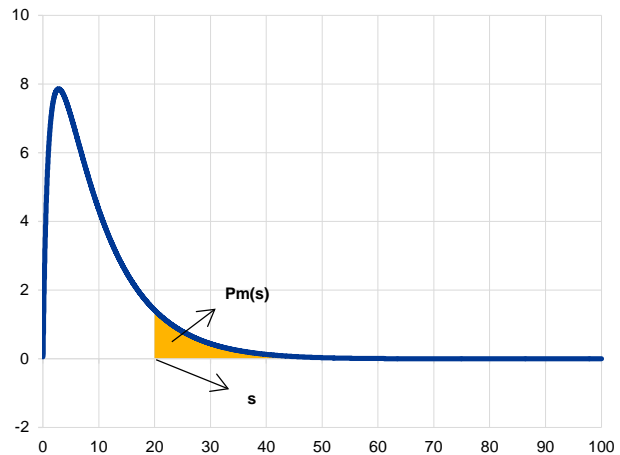
Measuring systemic risk

Probability density function

(x-axis: number of defaulting banks; percentages)
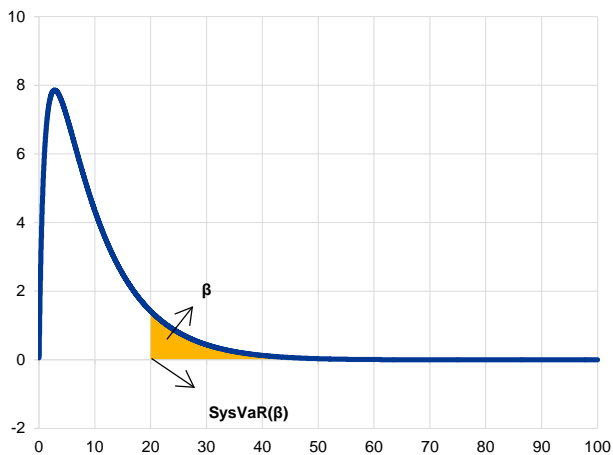
Probability density function

(x-axis: number of defaulting banks; percentages)

Source: ECB calculations.
Notes: The left-hand panel reports an example of the $SysVar(\beta)$. Given a confidence level β, $SysVar(\beta)$ represents the fraction of failing banks such that the probability of having more than $SysVar(\beta)$ defaults is equal to β. The right-hand panel reports the probability of having a systemic event. Given a threshold representing a number of bank defaults such that a systemic event occurs, $P_m(s)$ represents the probability of such an event occurring.

Box

Theoretical framework

This box shows that non-zero correlations between banks' default probabilities produce fat tails in the distribution of the number of bank failures. When such correlations are different from zero, the

probability that a given number of bank defaults is higher than a given threshold is larger than the case in which the same probability is computed under the assumption of zero correlations. As a consequence, looking at individual banks' default probabilities cannot guarantee financial stability.

Consider a set of N banks, indexed by $i = 1, 2, \ldots, N$. Each bank i is characterised by a given level of equity e_i. It is assumed that at any point in time $t \in [0, T]$ a bank defaults if its equity becomes negative. Regulators seek to ensure that the probability of banks' default remains below a given threshold by limiting the amount of risk which they can take on. Of course, regulators cannot reduce the probability of default to zero since this would mean that banks would not take any risk, including for example the risk deriving from lending to non-financial firms. For the banking system to play its role in the economy, the regulator has to tolerate that each bank bears a given risk of default. By imposing regulatory requirements, regulators decide the acceptable probability of failure α of each bank (in Basel II α is equal to 1/1000 and is defined over the time horizon (T) of one year).[174]

One can formalise these concepts as follows:

$$Pr\{e_i(t) \leq 0\} \leq \alpha, t \in [0, T], i = 1, 2, \ldots, N .$$

Case A. Banks' probabilities of default are uncorrelated

Under the assumption that banks' probabilities of default are uncorrelated, the banking system as a whole is relatively stable and the possibility that a financial meltdown occurs is remote. This can be easily shown by computing the probability of systemic events, i.e. the probability of a simultaneous default of a large number of banks over the time horizon [0, T]. To this end, let us consider the set of stochastic variables $\Theta_1, \ldots, \Theta_i, \ldots, \Theta_N$ which take on value one if bank i defaults, and value zero otherwise:[175]

$$\Theta_i(T) = \begin{cases} 1, if \ e_i(t) \leq 0 \text{ at any time } t \in [0, \text{T}] \\ 0, otherwise \end{cases} .$$

The total number of defaults over the considered time period is given by:

$$N_d(T) = \sum_{i=1}^{N} \Theta_i(T).$$

To compute the distribution of $N_d(T)$, let us exploit the fact that each variable $\Theta_i(T)$ follows a Bernoulli distribution:

$$\Theta_i(T) = \begin{cases} 1, with \ probability \ \alpha \\ 0, with \ probability \ (1 - \alpha) \end{cases} .$$

Since it is assumed that the pairwise default probabilities are uncorrelated – which is tantamount to assuming that the stochastic variables are independent – by the Central Limit Theorem (CLT), the distribution of $N_d(T)$ tends, for N large enough, to a Gaussian distribution characterised by mean $N\alpha$ and variance $N\alpha(1 - \alpha)$. When considering the percentage of the number of bank failures, the distribution of n_d, where $n_d = N_d/N$, will be:

---

[174] Note that this framework is a stylised version of the risk-based regulatory approach adopted, e.g. by the Basel Committee on Banking Supervision (BCBS). See Bank for International Settlements (2006), *Basel II: International Convergence of Capital Measurement and Capital Standards: A Revised Framework*, June. The BCBS assigns risk weights to assets to compute the amount of capital that banks have to hold. This enables banks to bear losses which could materialise for a given confidence level. This way, the regulator establishes the tolerable default probability for a bank.

[175] In this framework, we exclude the possibility of recoveries.

$$n_d(T) \sim g\left(\alpha, \frac{\alpha(1-\alpha)}{N}\right),$$

where $g(\cdot)$ denotes a Gaussian probability density function.

In practice this result holds when N is larger than 50. Having defined the distribution of $n_d(T)$, it is now possible to compute the probability that systemic events occur. In this framework systemic risk is defined as the likelihood that more than N_s defaults occur over the time period $[0,T]$. For very large N, defining $n_s = N_s/N$, the probability that n_d is greater than n_s is:

$$Pr\{n_d > n_s\} = \int_{n_s}^{1} g\left(x; \alpha, \frac{\alpha(1-\alpha)}{N}\right) dx.$$

To illustrate, let us assume that, in a banking system composed of $N = 1000$ banks, the regulator sets a threshold for the default probability equal to $\alpha = 0.001$. Then the probability of having a systemic event with $N_s = 20$ failures can be considered virtually inexistent.[176] In this framework, the soundness of each bank is enough to ensure financial stability.

Case B. Banks' probabilities of default are correlated

Consider now the case in which banks' default probabilities are correlated and let us explore the impact that such non-zero correlations have on systemic risk. To this end, it is necessary to compute the distribution of $N_d(T)$ when the assumptions for the CLT to hold are no longer valid. Following the approach of Vasicek (1987) for a loan portfolio,[177] now banks' equity levels are assumed to be correlated. This implies that the associated distribution of the percentage of the number of bank failures will be:

$$\tilde{n}_d(T) \sim \sqrt{\frac{1-\rho}{\rho}} \exp\left\{-\frac{1}{2\rho}\left(\sqrt{1-\rho}\,G^{-1}(n_d) - G^{-1}(\alpha)\right)^2 + \frac{1}{2}\left(G^{-1}(n_d)\right)^2\right\},$$

where $G(\cdot)$ denotes the cumulative Gaussian distribution function and ρ is the pairwise (non-zero) correlation among banks' default probabilities. This probability density function denotes the distribution of the number of defaults (expressed as a percentage of the total number of banks N) when the pairwise correlation between two banks' default probabilities is non-zero. Importantly, the default probability of each individual bank is still equal to α, as it was in the case where correlations were equal to zero. However, the distribution of the total number of defaults is different. This result also holds when the correlations between banks' default probabilities are not pairwise the same. Note that when $\rho = 0$, $\tilde{n}_d(T)$ collapses to $n_d(T)$.

By way of illustration, in line with the previous example, one can assume that the regulator tolerates a default probability equal to $\alpha = 0.001$ for each bank, and that the total number of banks is equal to $N=1000$. However, the correlation coefficient is now different from zero and equal to $\rho = 0.3$. In this case, the probability of having more than $N_s = 20$ defaults is roughly equal to 0.007: systemic events become plausible.

---

[176] Specifically, the probability of 20 banks failing is equal to 10^{-20}.

[177] See Vasicek O. (1987), *Probability of loss on loan portfolio*, KMV Corporation, February.

Data

In order to capture the overlapping portfolios in the euro area banking system, this article uses two relatively new ECB datasets. These datasets are the individual monetary financial institutions' (MFI) balance sheet data and the securities holdings statistics (SHS) data. Furthermore, the SHS dataset is complemented with information about capital and leverage ratios obtained from the European Banking Authority (EBA).

Owing to their relatively large time coverage, a subset of the individual MFI balance sheet data is used to compute the development of systemic risk over time (see the section entitled "Systemic risk in the euro area: the time dimension"). Data include observations from 2007 to 2014 at monthly frequency and cover around one hundred euro area MFIs. The MFI balance sheet data, however, do not provide detailed information on banks' overlapping portfolios.

To overcome this limitation, this special feature also makes use of SHS data, which contain granular security-by-security information on the overlapping portfolios of individual banks (see the section entitled "Systemic risk in the euro area: recent snapshots"). More specifically, the SHS dataset includes individual securities' holdings by the 26 largest banking groups headquartered in the euro area at quarterly frequency (SHS Group data).

To construct overlapping portfolios for the full euro area banking system, the SHS Group data are combined with SHS Sector data, which provide information on security-by-security holdings by the aggregate banking systems of the 19 euro area Member States.[178] Currently, the SHS dataset covers only a short time period – it is available as of the fourth quarter of 2013.[179]

Systemic risk in the euro area: the time dimension

By applying the methodology discussed in the previous section to the euro area banks covered in individual MFI balance sheet data, this section computes the dynamic evolution of the systemic Value-at-Risk when a contagion mechanism operates ($\mathrm{SysVaR}_t^c(\beta)$) and when it does not ($\mathrm{SysVaR}_t(\beta)$).[180] In particular, the systemic Value-at-Risk denotes the number of bank defaults (as a fraction of the total number of active banks) with a probability no larger than a given β. In this article β is set equal to 0.01.

When assuming no contagion, $\mathrm{SysVaR}_t(\beta)$ is constant over time and is equal to 2.8% (see the yellow line in the left-hand panel of Chart C.3). By contrast, the blue line reported in the same panel denotes the systemic Value-at-Risk when a contagion

---

[178] The security portfolios of the banks included in the SHS Group sample are subtracted from these banking system aggregates by country.

[179] For more information about SHS data, see ECB (2015), "Who holds what? New information on securities holdings", *Economic Bulletin*, Issue 2.

[180] The contagion mechanism is not operating if market liquidity is infinite, which implies that the price of securities does not change after a sale.

mechanism is operating. In particular, $\text{SysVaR}_t^c(\beta)$ represents, at each time t, the Value-at-Risk of the distribution of the number of bank defaults – i.e. the percentage of bank failures with a confidence level equal to $\beta = 0.01$. Unlike the $\text{SysVaR}_t(\beta)$, $\text{SysVaR}_t^c(\beta)$ varies over time, since its computation takes into account a time-varying deleveraging process which reflects variations in banks' balance sheets. The distance between the blue line and the yellow line denotes the systemic risk which derives from contagion and, in particular, fire sales.

At the beginning of the sample, the systemic Value-at-Risk is computed under the assumption that there is contagion as high as 10%, which means that, with a probability of 1%, more than 10% of the banks in the sample could fail. The $\text{SysVaR}_t^c(\beta)$ reaches its peak at the end of 2008, when more than 13% of the banks could go bust with a probability of 1%, to decline sharply thereafter. In the last part of our sample, the blue line and yellow line coincide, which implies that the contagion mechanism is not playing any role. Changes in the level of systemic risk can be due to changes in the banks' security portfolios or changes in banks' capital.

The left-hand panel of Chart C.3 provides an important policy message: when considering idiosyncratic default probabilities, it is necessary to take banks' interconnections into account in order to capture how a bank idiosyncratic shock can reverberate across the whole banking system and become systemic. Moreover, since the Value-at-Risk is computed at a relatively high confidence level (1%), and since there is an upper bound for the banks' default probabilities, under the assumption of no contagion, ensuring the stability of individual banks would be sufficient to guarantee the stability of the whole system – but since banks' default probabilities are positively correlated such an approach is, in fact, insufficient to preserve stability in the system as a whole.[181]

This framework can also be used to compute the probability that a systemic financial crisis occurs. By setting the percentage of banks going bust simultaneously – which here is set at 5% – it is possible to estimate the probability that such a systemic event occurs.[182] Such probability, which is depicted by the blue line in the right-hand panel of Chart C.3, increases sharply in the second half of 2007, reaching its peak in March 2008. As of 2010, this probability becomes negligible and it is indistinguishable from the probability of a systemic event when there is no contagion

---

[181] In line with the regulatory framework, we assume that the idiosyncratic individual bank default probability is equal to 0.001, which is an upper bound. In principle, one should use the real banks' default probabilities. However, since such probabilities are pro-cyclical, it is preferred to keep them constant at their upper bound and study how variations in balance sheets affect the probability of systemic events. This allows us to isolate a particular contagion mechanism – the fire sales – from other factors which could influence systemic risk measures.

[182] Although the number of yearly bank defaults in the United States from 1934 to 2014 suggests that a deep financial crisis ensues when a fraction equal to or larger than 3% of the banking system fails simultaneously, the sample under consideration includes a relatively small number of banks (roughly one hundred). Therefore, a conservative approach is adopted, defining systemic events as those characterised by at least 5% of simultaneous bank defaults.

in the banking system. In the case no contagion takes place (yellow line), the probability that a systemic events occurs is constant and equal to zero.[183]

Chart C.3

Systemic risk: the time dimension

Systemic VaR	Probability of systemic event
(percentages)	(percentages)
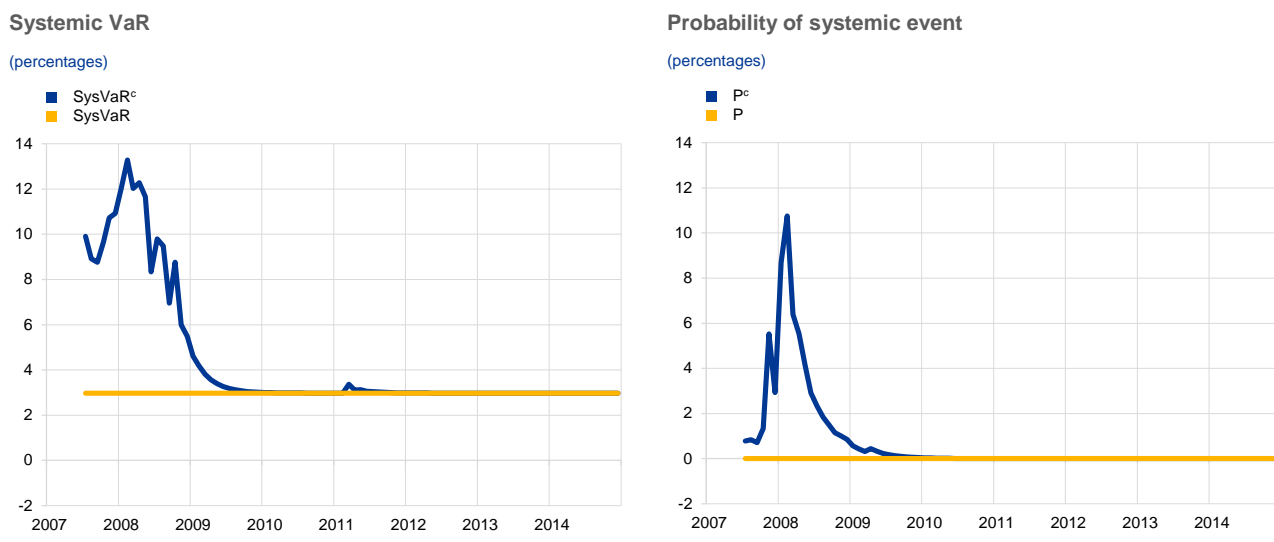
Sources: ECB (individual MFI balance sheet items statistics) and ECB calculations.
Note: The left-hand panel reports the $SysVar(1\%)$ and the right-hand panel represents the probability of having a systemic event $P_m(5\%)$, both in the case where contagion occurs (blue line) and in the case there is no contagion (yellow line).

Systemic risk in the euro area: recent snapshots

In this section, the systemic Value-at-Risk is computed using a different dataset, the securities holdings statistics (SHS). Although the time length of this dataset is rather short – observations start in the fourth quarter of 2013 – its fine granularity enables us to obtain recent snapshots of systemic risk estimates, which account for the network of securities' overlapping portfolios.

The left-hand panel of Chart C.4 reports the systemic Value-at-Risk with a confidence level equal to 1% computed in two cases, i.e. the case in which a contagion mechanism operates, and the case in which no fire sales occur. Although the marginal default probabilities of banks are the same in the two cases, correlations induced by common exposures increase the fragility of the financial system. The right-hand panel of Chart C.4 reports the probability that a systemic

---

[183]   The sovereign debt crisis is not captured by the measures of systemic risk proposed here. The reason is that these measures do not consider any price shock which is not generated by the financial system itself. After shocking the system by letting banks fail according to a probability specified by the regulator (1/1000), fire sales are triggered. The subsequent systemic Value-at-Risk only captures the amount of systemic risk attributable to fire sales, but not to other shocks such as the decline in the sovereign debt value. However, the framework is sufficiently general to accommodate further sources of shocks.

event occurs, defined as the failure of a fraction of the banking system larger than 5%.[184]

Under the assumption that contagion takes place, the systemic Value-at-Risk and the probability that a systemic crisis occurs contract significantly in the first quarter of 2014. This decrease is likely related to the announcement of the comprehensive assessment on 23 October 2013.[185] After the announcement, banks increased their capital and reshuffled their security portfolios. This contributed to reducing systemic risk.
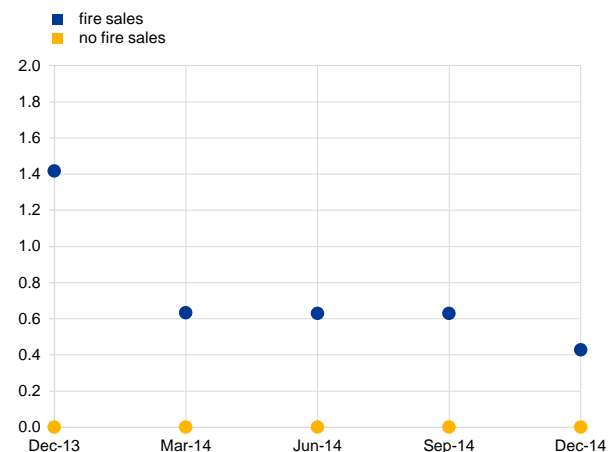
Chart C.4

Systemic risk: Q4 2013 – Q4 2014

Systemic VaR	Probability of systemic event
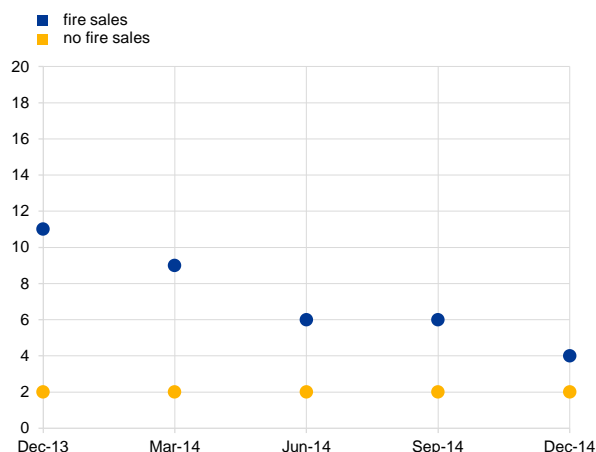
(percentages) · (percentages)

Sources: ECB (SHS Group and SHS Sector), European Banking Authority, and ECB calculations.
Note: The left-hand panel reports the $SysVar(1\%)$ and the right-hand panel represents the probability of having a systemic event $P_m(5\%)$, both in the case where contagion occurs (blue dots) and in the case there is no contagion (yellow dots).

Finally, the left-hand panels of Charts C.5 and C.6 report the networks of overlapping portfolios. Each node represents a bank in the sample. Two nodes are connected if there is an overlap in the banks' tradable securities portfolios. Colour and size of the nodes highlight their centrality. By the same token, the colour and thickness of the links highlight how large the common exposure is. These charts illustrate how the topology of the banks' network changed after the announcement of the comprehensive assessment. The right-hand panels of Charts C.5 and C.6 instead report the distributions of the number of bank defaults in the fourth quarter of 2013 and in the fourth quarter of 2014, which also changed after the announcement of the comprehensive assessment. As a consequence, the systemic Value-at-Risk

---

[184] In this sample of 45 banks, 5% of failures correspond to roughly three banks going bust. With a larger sample this measure would produce a more realistic number of failures. However, since the article considers the 26 largest euro area banks and 19 bank aggregates by country, when more than three entities fail, this can certainly be associated with a systemic event.

[185] The impact of other events cannot be ruled out. For instance, the European Banking Authority (EBA) published the results of the transparency exercise in December 2013 and banks have also improved their capital levels in advance of changes to the regulatory framework.

computed at a confidence level of 1% decreased from 11% in 2013 Q4 to 4% in 2014 Q4 (see the yellow areas).
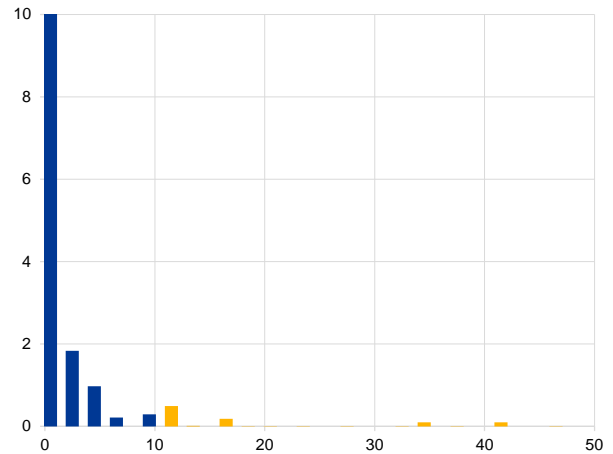
Chart C.5

Systemic risk: Q4 2013

Network of overlapping portfolios
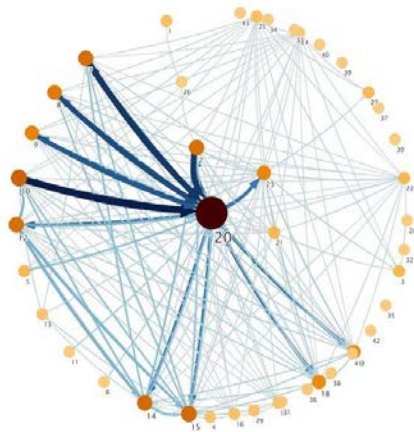
Simulated distribution of number of bank defaults

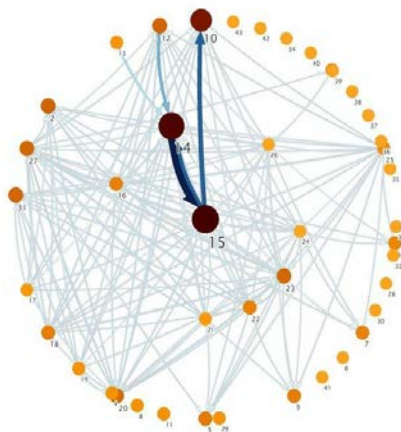(x-axis: number of defaulting banks; percentages)

Sources: ECB (SHS Group and SHS Sector), European Banking Authority, and ECB calculations.
Notes: The left-hand panel reports the network of overlapping portfolios. Each node represents a bank in the sample or a banking system aggregate for a country. Two nodes are connected if there is an overlap in the banks' tradable securities portfolios. Colour and size of the nodes highlight their centrality. Colour and thickness of the links highlight how large the common exposure is. The right-hand panel reports the simulated distribution of the number of bank defaults in 2013 Q4. The yellow areas show the 1% quantile of the distribution.
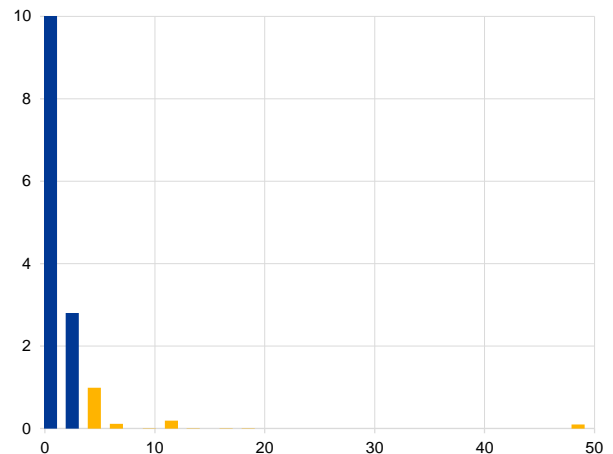
Chart C.6

Systemic risk: Q4 2014

Network of overlapping portfolios

Simulated distribution of number of bank defaults

(x-axis: number of defaulting banks; percentages)

Sources: ECB (SHS Group and SHS Sector), European Banking Authority, and ECB calculations.
Notes: The left-hand panel reports the network of overlapping portfolios. Each node represents a bank in the sample or a banking system aggregate for a country. Two nodes are connected if there is an overlap in the banks' tradable securities portfolios. Colour and size of the nodes highlight their centrality. Colour and thickness of the links highlight how large the common exposure is. The right-hand panel reports the simulated distribution of the number of bank defaults in 2014 Q4. The yellow areas show the 1% quantile of the distribution.

Concluding remarks

This special feature defines systemic risk as the simultaneous failure of a "large" number of banks. The notion of "large" is qualified by looking at disruptive financial crises in the United States from 1934 to 2014. In such crises more than 3% of banks defaulted at the same time. Exploiting this intuition, this article suggests a measure of systemic risk as the Value-at-Risk of a banking system, i.e. the percentage of banks going bust simultaneously over a given time horizon for a given confidence level. To estimate the systemic Value-at-Risk, the distribution of the number of bank failures is derived. In this framework, the mechanism generating fat tails in such a distribution and therefore leading to systemic risk is contagion. In particular, contagion materialises through fire sales and is affected by the topology of the network of banks' common exposures. The framework is general enough to accommodate any contagion mechanism.

This special feature applies this framework to data on the euro area banking system. After the announcement of the comprehensive assessment in October 2013 banks reshuffled their security portfolios, which resulted in a decline in the probability of a systemic event occurring.

The framework proposed in this special feature has significant policy implications. In contrast to the monetary policy domain where extensive literature exists on the definition and measurement of price stability, no equivalent, quantifiable objective is available to macroprudential policy-makers. This special feature seeks to fill this gap. A clear definition and measurement of systemic risk can enhance the design of policies to contain it and contribute to the accountability of policy-makers.

The framework can also be extended to identify systemically important assets and banks and to track their systemicness over time. It therefore allows policy-makers to take appropriate measures to reduce the likelihood that systemic risk materialises and target the main factors responsible for driving systemic events.