The Price of Complexity in Financial Networks

Stefano Battiston [*], Guido Caldarelli [†] [‡] [§], Robert M. May [¶], Tarik Roukny [‖] [**], Joseph E. Stiglitz [††]

[*]Dept. Banking and Finance, University of Zurich, Switzerland,[†]IMT Institute for Advanced Studies, Piazza San Francesco 19, 55100 Lucca, Italy,[‡]ISC-CNR, Via dei Taurini 19, 00185 Roma, Italy,[§]London Institute for Mathematical Science, South St. 35 Mayfair, London W1K 2XF, UK,[¶]Oxford University, Oxford OX1 1HP, UK,[‖]IRIDIA, École Polytechnique & Solvay Brussels School of Economics and Management, ULB, Belgium,[**]Fonds National de la Recherche Scientifique, Belgium, and [††]Columbia University, New York 10027, USA

Submitted to Proceedings of the National Academy of Sciences of the United States of America

Financial institutions form multi-layer networks of contracts among each other and exposures to common assets. As a result, the default probability of one institution depends on the default probability of all the other institutions in the network. Here, we show how small errors on the knowledge of the network of contracts can lead to large errors on the probability of systemic defaults. From the point of view of financial regulators, our findings show that the complexity of financial networks may decrease our ability to mitigate systemic risk and thus it may increase the social cost of financial crises.

financial markets | systemic risk | complex systems | computational social science

Several years after the beginning of the so-called Great Recession, regulators warn that we still do not have a satisfactory framework to deal with too-big-to-fail institutions and with systemic events of distress in the financial system[1, 2]. The topic is of general societal and scientific interest since assisting financial institutions in the recent years has come with a social cost equivalent to few percent points of GDP for OECD countries [3]. In particular, poor estimates of systemic risk are socially costly because regulators and banks' managers end up keeping aside either insufficient or redundant buffers [22, 23, 24, 25]. One of the difficulties is that financial institutions are connected in multi-layer networks both directly via contracts among each other (e.g. loans, bonds, repurchasing agreement, derivatives, etc.), and indirectly via exposures to common assets [4, 5, 6, 7, 8, 9, 10, 11]. The default probability of one institution depends therefore on the default probability of all the other institutions in the network. The resulting complexity of the financial system is a potential source of information asymmetries, collective moral hazard and increased systemic risk [18, 19, 20, 21], and hence requires deeper understanding. In particular, the determination of the probability of systemic events has remained an open problem so far [12, 13, 14, 15, 16, 17]. Here, we show that in a network of financial contracts the probability of systemic default can be very sensitive on errors on the information about contracts. Moreover, the sensitivity depends on the network structure due to a multiplicative interplay of errors along chains of lending. Under certain conditions, network structures with more numerous and longer chains of lending lead to a stronger amplification of errors. While there may be some intuition for this effect in the literature, it had not been explained analytically or quantified before. Our work helps to understand how to contain errors on the estimation of the probability of systemic events.

Results

We introduce a general model of a network of credit contracts among n financial institutions (hereafter, "banks"). Contracts are over-the-counter, i.e. they are not mediated through a central counterparty. They are also collateralized, i.e. banks have to post a collateral in order to receive a loan. In addition to making contracts with each other, banks also hold external assets, i.e. securities that are not issued by the banks in the system. These external assets are the only source of stochasticity in the model. An illustration of the case of three banks connected through credit contracts is sketched in Figure 1 (see the Methods section for the details of the model).

Errors on Contract Characteristics. We consider a first scenario, named here as "errors on contract characteristics", in which it is known which assets and which counterparties each banks is exposed to, but there can be an error, for instance on the recovery rate R (i.e. the fraction of the face value of the loan that can be recovered after the default of a counterparty). We first study how the systemic default probability P^{sys} varies as a function of parameter errors in the simplest case of two symmetric banks. We explore all pairs of parameter values in a range consistent with empirical evidence on interbank markets [4]. In Figure 2, each pair of curves with the same color represents the maximum and minimum value of the default probability as a function of the deviation on each parameter around a given point. Figure 2 (left) refers to the case in which the actual value (i.e. with zero error) of the default probability is 1. This means that for instance with an error on recovery rate R of 20%, one may think that the default probability is 0.4 when it is actually 1. Conversely, Figure 2 (right) refers to the case in which the actual value of the default probability is 0. For instance, with an error on the expected asset return μ of 20%, one may think that the default probability is 1 when it is actually 0. In both cases, the gap between the possible estimates increases rapidly with the error on the parameters, e.g. an error on parameters of 10% can lead to an error of 100% on the probability. Note that this result does not imply that the error on the default probability is *always* large compared to errors on the parameters. However, it shows that there exist cases in which the deviation can be very large. Next, in order to illustrate how the

Significance

Estimating systemic risk in networks of financial institutions represents today a major challenge both in science and financial policy making. This work shows how the increasing complexity of the network of contracts among institutions comes with the price of increasing inaccuracy in the estimation of systemic risk. The paper offers a quantitative method to estimate systemic risk and its accuracy.

Reserved for Publication Footnotes

effects of errors on contract characteristic depend also on the structure of the underlying network of contracts, we focus on errors on a single parameter, namely the recovery rate R, and we consider three basic architectures with three nodes: a star, a chain and a ring. Figure 3 shows the sensitivity $\partial P^{\text{sys}}/\partial R$ of the default probability on the recovery rate R as a function of the ratio between interbank leverage β and the maximal loss $\epsilon\sigma$ that a bank can withstand on the external assets. As we can see, the sensitivity is highest for the ring network architecture, followed by the chain and the star architectures. The intuition behind this result is that the systemic default probability in an interbank network depends on the multiplicative interplay among the parameters that matter for the default thresholds. In particular, the multiplication involves the banks located along chains of lending. As a result, more numerous and longer chains lead to stronger amplification of the errors on the parameters. Indeed, as proved in Methods, $\partial P^{\text{sys}}/\partial R$ is a polynomial in the interbank leverage β, where the leading power depends on the presence of chains or cycles in the network. This result, which had not been reported so far in the literature, illustrates concretely the impact of financial complexity on the determination of systemic risk.

Errors on the Structure of the Contract Network. Finally, we would like to investigate how the systemic default probability depends on the complexity of the network of contracts. To this end, we consider a second scenario, named here as "errors on network structure", in which the information regarding how many contracts a bank has and with which counterparties may be incorrect. More precisely, we are interested in measuring the error on the systemic default probability when the arrangement of contracts (i.e. who trades with whom) is not known and the number of possible arrangements increases, subject to the constraint of a maximum number of possible contracts, referred to as "density cap" (see Methods). In Figure 4, the green area corresponds to a lower expected return on external assets μ and the blue area to a higher one. Both areas represent the range of possible values of the default probability resulting from all possible network configurations compatible with a given density cap. The tip of each arrow indicates the value of default probability for three specific network architectures, i.e. the *star*, the *circle* and the *complete* graph. For instance, in the case of lower μ (green area), the complete graph yields values of default probability larger than the star architecture. In this context, a simple way to capture the complexity of the network is to count the number of possible network configurations for a given density cap. In the same figure, the red curve represents the entropy $\Sigma(C_l)$ of the space of possible network configurations (see Methods) for increasing values of the density cap.

Discussion

The complexity of the financial system is a source of potential information asymmetries and collective moral hazard [18, 26]. Indeed, if the default of some financial institutions has an impact on the system that is possibly very large but difficult to compute exactly, those institutions are more likely to enjoy a bail-out with public funds. They tend therefore to count on being rescued in case of downturn and take more risk than they would otherwise. In other words, such moral hazard leads to a more fragile financial system and to an implicit public subsidy. Therefore, the estimation of the probability of individual and systemic events in a network of contracts is crucial to improve the stability of the financial system. Yet the interdependence among asset values and probability of default of all institutions poses conceptual and computational challenges and little progress has been made in this direction so far. Many previous works on systemic risk build on the approach *à la* Eisenberg and Noe [27, 28, 4, 29, 30] in which a clearing vector of payment and a recovery rate on defaulting banks' assets is determined endogenously. However, the reason why the recovery rate can be determined endogenously is that the valuation is carried out *ex-post* (i.e. at the maturity of the contracts) and that, in case of default, there has been a successful and immediate asset liquidation of the external assets with full recovery. In contrast, in the practice, asset liquidation implies legal settlements that take several months or years, while in the short run (e.g. weeks) the recovery rate is likely to be significantly smaller than in the Eisenberg-Noe approach and to be sensitive to regulators' interventions. For these reasons, it is important to study the impact of errors on the recovery rate. Moreover, and even more importantly, here we are interested in the valuation that can be carried out *ex-ante*, i.e. before the maturity, given the information available. Therefore, we assume that banks lend to each other against a collateral and that the recovery rate on the loans to defaulted banks is smaller than one and exogenous [17]. Finally, similar to [5, 20, 21], we intentionally leave aside at this stage the question of which configurations and parameters would arise from banks' dynamic choice, since our method encompasses all possible configurations including those arising from the uncoordinated individual investment strategies.

On the one hand, the ability of a bank to make contracts with any other bank in the system increases its ability to diversify the risk. On the other hand, the resulting complexity comes with the price that "everybody knows less". Indeed, both a higher complexity of the individual contracts and a higher complexity of the structure of contracts implies that market participants and regulators know less precisely the probability of individual and systemic default. While there are individual incentives to be part of a complex financial network, this work shows quantitatively the existence of so-called "negative externalities", which eventually translate into potential social costs.

More generally, our results show that higher interdependence on the credit market among banks decreases our ability to make predictions on defaults. Hence a potential tradeoff emerges between financial stability and market complexity. However, market complexity may be poorly addressed by complex regulations, due to the intrinsic limitations related to the precision and accuracy of the estimations of the probability of default, as highlighted in this work. Our results corroborate the view of a regulatory response grounded in simplicity [31].

Methods

We consider a financial network with over-the-counter (OTC) contracts among n banks, including secured credit contracts. We distinguish between contracts within the banking system itself ("interbank") and contracts of banks on securities outside the banking system ("external").

Model timing. The timing of the model is as follows. At time 1, banks raise funds and make investments in external and interbank assets. At time 2, the values of the external assets are shocked and updated. While the shock distribution is known at time 1, shocks are only observed at time 2. At time 2 the interbank contracts mature and their value is also updated depending on the shocks that have occurred. Figure 1 sketches the structure of our model. For each bank i, the main quantities are the following.

Assets and liabilities. Assets and liabilities of i on the external markets are denoted as a_i^E and ℓ_i^E. Assets and liabilities of i on the interbank credit market are denoted as a_i^B

and ℓ_i^B. Total liabilities are denoted as ℓ_i. At time 1, each bank i allocates its external assets in a portfolio of securities on the external market, E_{ik} denoting the fraction of i's external assets invested at time 1 in the security k. The unitary value of the external security k is x_k^E. Without loss of generality: at time 1, $x_k^E(1) = 1$ for all k, while $x_K^E(2)$ is a random variable drawn from a given distribution. At time 2, then the external assets of bank i, is a sum of random variables, $a_i^E(2) = a_i^E(1)\sum_k E_{ik}x_K^E(2)$. For our purposes, it is sufficient to assume that we can express the external assets of bank i as follows: $a_i^E(2) = a_i^E(1)(1 + \mu + \sigma\, u_i)$, where u_i is a random variable drawn from a given distribution with mean zero and variance one, the parameter μ_i is the expected return of the portfolio and σ_i its standard deviation. We also assume to know the joint probability distribution $p(u_1, ..., u_n)$.

At time 1, each bank i allocates its interbank assets among the other banks, B_{ij} denoting the fraction of i's interbank assets invested at time 1 in the liability of bank j. These investments are secured via collateralisation, i.e. bank j posts as collateral for the loan an asset that bank i will collect in case bank j defaults. In spirit, this approach is similar to [21], although for the purpose of our study, we exclude rehypothetication, i.e. assets used as collateral are kept aside and cannot be re-used. The collateral is also assumed to be risk-free in the time-horizon of the model, i.e. its value stays constant during the two periods and it can be ignored in the default condition further below.

The unitary value of the interbank liability of bank j to other banks is x_j^B. Without loss of generality: at time 1, $x_j^B(1) = 1$ for all j. The liabilities of bank j are constant in value from the perspective of bank i, i.e. the debt agreed upon in the contract at time 1. However, from the point of view of counterparties of j, $x_j^B(2) = 1$ if bank j honors its obligation, $x_j^B(2) = R_j$ otherwise, where R_j is the recovery rate, i.e. the fraction of the interbank asset that is covered by the collateral and that the lender can recover after the default of j. Accordingly, at time 2, the interbank assets of bank i, is $a_i^B(2) = a_i^B(1)\sum_j B_{ij}x_j^B(2)$. Notice that this approach differs from previous models [27, 29] based on the determination ex-post of the clearing vector of payments.

Default condition. The standard balance-sheet identity in financial accounting states that equity of bank i, e_i, is the difference between assets and liabilities. Hence $e_i(2) = a_i^E(2) + a_i^B(2) - \ell_i = a_i^E(1)(1+\mu+\sigma\, u_i) + a_i^B(1)\sum_j B_{ij}x_j^B(2) - \ell_i$. It is also standard to assume that the default of bank i occurs when equity becomes negative, i.e. if $e_i(2) < 0$. In the following, we are interested in the probability of default of individual banks. Notice that we assume $e_i(1) > 0$, thus $e_i(2) < 0$ iff $\frac{e_i(2)}{e_i(1)} < 0$. It is then convenient to write the default condition as $\varepsilon_i(1 + \mu + \sigma\, u_i) + \beta_i \sum_j B_{ij}x_j^B(2) - \lambda_i < 0$, where the parameter $\varepsilon_i = \frac{a_i^E(1)}{e_i(1)}$ measures the magnitude, per unit of initial equity of bank i, of the investments of bank i in external assets. Similarly, the parameter $\beta_i = \frac{a_i^B(1)}{e_i(1)}$ measures the magnitude, per unit of initial equity, of i's investments in interbank assets and the parameter $\lambda_i = \frac{l_i(1)}{e_i(1)}$ measures the magnitude, per unit of initial equity, of i's total liabilities. Let us define a default indicator χ_i, with $\chi_i = 1$ in case of default of bank i and $\chi_i = 0$ otherwise. Because the only variable that is exogenously stochastic is the shock u_i on each bank external assets, we finally write the default condition as follows

$$u_i < \theta_i \equiv \frac{1}{\varepsilon_i \sigma}(\lambda_i - 1 - \mu_i - \beta_i \sum_j B_{ij}x_j^B(\chi_j)), \qquad [1]$$

where θ_i denotes the default threshold. Notice the we have dropped the time in the notation, while we have emphasized in the formula that the value of the interbank liability of a counterparty j, x_j^B, depends on the default indicator of j, χ_j, to recall that it is $x_j^B(\chi_j = 0) = 1$ and $x_j^B(\chi_j = 1) = R_j$. Thus, depending on the magnitude and sign of the shocks u_i that hit all banks, some of them can default on their obligations, possibly causing other banks to default. We can now express the default indicators χ_i of all banks as a system of equations

$$\forall i \quad \chi_i = \Theta(u_i - \theta_i(\chi_1, ..., \chi_n)), \qquad [2]$$

where Θ denotes the step function or Heaviside function (i.e. equal one if the argument is positive, zero otherwise). A solution of the system, denoted as χ^*, depends on the vector of shocks u and on the initial condition χ^0, which represents the initial belief of the banks in other banks' default. The existence and uniqueness of the solution is discussed further below.

The determination of the fixed point of the map above can become computationally cumbersome if we want to sample at a fine resolution the shock space of an arbitrary number of banks. However, in the default condition of the counterparties j of bank i, we can simplify the computation by replacing the value of the second-order counterparties' credit obligations (i.e. the obligations of neighbors of order 2 in the contract network) with their expected value. In other words, we replace the stochastic variable x_k^B with its expected value $E[x_k] = RP_k + (1 - P_k)$, where P_k is the default probability of bank k. This is a legitimate approximation since expected values are commonly used in the banking practice to estimate the future value of assets. Notice also that this approximation does not remove the effects of correlations across shocks on banks.

Default probability. The probability P_i of default at time 2 of bank i is then simply the integral over the shock space of the default indicator:

$$\forall i \quad P_i = \int \chi_i^*(u, \chi^0) p(u) du, \qquad [3]$$

where $p(u)$ denotes the joint density function of the shocks and accounts for possible correlations across shocks. Finally the probability of systemic default is

$$P^{sys} = \int \chi^{sys}(\chi^*(u, \chi^0)) p(u) du, \qquad [4]$$

Where χ^{sys} is the systemic default indicator. The choice of the systemic default identification can vary and there is no consensus on what should be defined as a systemic event. For the sake of clarity, in the following we consider the extreme but intuitive case of all banks defaulting, i.e. $\chi^{sys} = \Pi_i \chi_i^*$.

Existence and uniqueness of solutions. In our model, banks's default conditions are described by a system of non-linear equations, $\chi = \Theta(u, \chi)$, for which no closed-form solution exists. However, finding all the solutions of such system is equivalent to study, for any point u of the shock space and for any initial condition on χ, the deterministic map of the finite set $\{0, 1\}^n$ in itself: $\Theta : \{0, 1\}^n \to \{0, 1\}^n$. Because the map is deterministic and the set is finite, for any given initial condition it must have either a unique fixed-point or a limit cycle, otherwise the set would have to be non-finite [32]. If banks only make regular credit contracts with each other, then there is a unique solution. Indeed, the default state of a bank can only be affected adversely by the default of some counterparty. Hence the vector of default states of the banks is also a non-decreasing function of the default state of the others. This implies that the map cannot enter a cycle involving more than one state. Otherwise the default state of at least one bank would have to revert from default to no-default, which is not possible by assumption. Hence, the fixed-point is unique for any given initial condition on χ.

Contract Errors. By contract errors we mean that the estimates of the contract characteristics available to the regulator who computes the systemic default probability deviate from the actual value. We define the relative error on a given parameter p, the ratio $\Delta_p = \frac{\tilde{p}-p^*}{p^*}$, i.e. the relative deviation of the regulator's estimation, $\tilde{p}$, from the actual value p^* of the variable p. For a given level of deviation Δ^*_p, we consider a realistic number giving a Δ_p precision of 5% around the original value p^*: $\Delta_p \in [p^* - \Delta^*_p, p^* + \Delta^*_p]$. For instance, a deviation of 10% on the the interbank leverage β (i.e. $\Delta^*_\beta = 0.1$), will yield $\Delta_\beta \in \{-10\%, -9\%, ..., 0, ..., 10\%\}$. For each value of Δ_β, we compute the corresponding probability of systemic default using $\tilde{\beta} = \beta^*(1 + \Delta_\beta)$ and we record the minimum and the maximum values. Additionally, we consider the effect of combined errors where more than one parameter is subject to changes at the same time. Take as example the leverage parameters ε and β. As $\tilde{\varepsilon}$ and $\tilde{\beta}$ can each take 21 different values, we have 21^2 combinations. In the spirit of what has been done previously, we compute all probabilities of systemic default using all possible combinations of $\tilde{\varepsilon}$ and $\tilde{\beta}$ given a level of deviation $\Delta^*_{(\varepsilon,\beta)}$.

Analytical example. In order to understand how the uncertainty of specific parameters can be amplified by the network structure, consider the following three basic architectures with three nodes: a star, a chain and a ring. Here we compute the probability P^{sys} of the event in which all banks default under the assumption that the shocks hitting the banks are independent and drawn from the same uniform distribution in the space $[-1, 1]$, i.e. $p(u_i) = 1$, and $p(u_i, u_j) = p(u_i)\, p(u_j)\, \forall i, j$. Denote by θ_i^+ the specific value of the threshold θ_i in Eq. 1 in the case that all of the counterparties of a given bank i default. Conversely, denote by θ_i^- the case of no defaulting counterparties. In particular, if a bank has no counterparties (e.g. a leaf node in a tree network structure) then it has $\theta_i = \theta_i^-$. The computation yields $P^{\text{sys,star}} = (1/2^3)(1+\theta_1^+)(1+\theta_2^-)(1+\theta_3^-)$; $P^{\text{sys,chain}} = (1/2^3)(1+\theta_1^+)(1+\theta_2^+)(1+\theta_3^-)$; $P^{\text{sys,ring}} = (1/2^3)(1+\theta_1^+)(1+\theta_2^+)(1+\theta_3^+)$. Note that $\theta_i^+ \geq \theta_i^-$ because a bank needs more positive shocks on its external assets in order to survive when counterparties on the interbank have defaulted. It follows that $P^{\text{sys,ring}} \geq P^{\text{sys,chain}} \geq P^{\text{sys,star}}$. It also follows that the sensitivity of the default probability on the recovery rate R depends on the network structure: $\partial P^{\text{sys,ring}}/\partial R = (\beta/(\epsilon\,\sigma))^3$; $\partial P^{\text{sys,chain}}/\partial R = (\beta/(\epsilon\,\sigma))^2$; $\partial P^{\text{sys,star}}/\partial R = \beta/(\epsilon\,\sigma)$. In this example, as long as $\beta/(\epsilon\,\sigma) > 1$, which is empirically plausible in many cases, the sensitivity on errors on the recovery rate can be much larger in the ring than in the star.

Network Errors. By network errors we mean that the regulator does not have information on the arrangement of the contracts among banks but only their maximum number. We denote by C_l the cap on the maximum number of possible contracts in the market (i.e. the network density). By increasing the cap C_l, the total number of possible configurations grows as follows: $\overline{\mathcal{N}}(C_l) = \sum_{l=0}^{C_l} \binom{n(n-1)}{l}$. For a given C_l, we inspect all possible configurations of links arrangement and compute the probability of systemic default. As a measure of the market complexity we take the so-called network entropy, relative to a given density cap [33], i.e. the logarithm of the number of configurations normalized by the number of nodes, $\Sigma(C_l) = (1/n)\, log\overline{\mathcal{N}}(C_l)$.

Available Code and Data. Analyses and figures can be reproduced using Python, Matlab and C++ scripts and data files available in the public GitHub repository following the instructions provided in the repository.

ACKNOWLEDGMENTS. SB and GC acknowledge support from: FET Project SIMPOL nr. 610704, FET project DOLFINS nr 640772, and FET IP Project MULTIPLEX nr 317532. SB acknowledges the Swiss National Fund Professorship grant no. PP00P1-144689. TR acknowledges the support of Fonds National de la Recherche Scientifique, FNRS, Belgium. JES and SB acknowledge the INET grant on Financial Stability. All the authors developed the concepts of the study. SB, GC and TR conducted the analysis, and the statistical methods needed for data interpretation. All the authors wrote the manuscript.

1. Haldane, A. G. & May, R. M. Systemic risk in banking ecosystems. Nature 469, 351–5 (2011).
2. BoE. A framework for stress testing the UK banking system. Tech. Rep. October (2013).
3. Lambert, F. J., Ueda, K., Deb, P., Gray, D. F. & Grippa, P. How big Is the implicit subsidy for banks considered too important to fail? Global Financial Stability Report, April 2014, SSRN 2419118 (2014).
4. Elsinger, H., Lehar, A. & Summer, M. Risk Assessment for Banking Systems. Management Science 52, 1301–1314 (2006).
5. Beale, N. et al. Individual versus systemic risk and the Regulator's Dilemma. Proceedings of the National Academy of Sciences 108, 12647–12652 (2011).
6. May, R. M. & Arinaminpathy, N. Systemic risk: the dynamics of model banking systems. Journal of the Royal Society, Interface / the Royal Society 7, 823–838 (2010). URL .
7. Anand, K., Gai, P. & Marsili, M. Rollover risk, network structure and systemic financial crises. Journal of Economic Dynamics and Control 36, 1088–1100 (2012).
8. Acemoglu, D., Ozdaglar, A. & Tahbaz-Salehi, A. Systemic Risk and Stability in Financial Networks. American Economic Review 105, 564–608 (2015).
9. Elliott, M., Golub, B. & Jackson, M. O. Financial Networks and Contagion. American Economic Review 104, 3115–3153 (2014).
10. Peltonen, T. A., Scheicher, M. & Vuillemey, G. The network structure of the CDS market and its determinants. Journal of Financial Stability 13, 118–133 (2014).
11. Duffie, D., Scheicher, M. & Vuillemey, G. Central clearing and collateral demand. Journal of Financial Economics (2015).
12. Stiglitz, J. E. Risk and Global Economic Architecture: Why Full Financial Integration May Be Undesirable. American Economic Review 100, 388–392 (2008).
13. Gai, P. & Kapadia, S. Contagion in financial networks. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 466, 2401–2423 (2010). URL .
14. Markose, S., Giansante, S. & Shaghaghi, A. R. Too Interconnected To FailFinancial Network of US CDS Market: Topological Fragility and Systemic Risk. Journal of Economic Behavior & Organization 83, 627–646 (2012).
15. Battiston, S., Delli Gatti, D., Gallegati, M., Greenwals, B. & Stiglitz, J. E. Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk. Journal of Economic Dynamics and Control 36, 1121–1141 (2012).
16. Battiston, S., Puliga, M., Kaushik, R., Tasca, P. & Caldarelli, G. DebtRank: Too Central to Fail? Financial Networks, the FED and Systemic Risk. Scientific Reports 2, 1–6 (2012).
17. Cont, R., Moussa, A., Santos, E. B. & Others. Network structure and systemic risk in banking systems. Handbook of Systemic Risk 327–368 (2013).
18. Battiston, S., Caldarelli, G., Georg, C.-P., May, R. & Stiglitz, J. Complex derivatives. Nature Physics 9, 123–125 (2013).
19. Caballero, R. J. & Simsek, A. Fire sales in a model of complexity. The Journal of Finance 68, 2549–2587 (2013).
20. Arinaminpathy, N., Kapadia, S. & May, R. M. Size and complexity in model financial systems. Proceedings of the National Academy of Sciences 109, 18338–18343 (2012).
21. Gai, P., Haldane, A. & Kapadia, S. Complexity, concentration and contagion. Journal of Monetary Economics 58, 453–470 (2011).
22. Rochet, J. C. Why are there so many banking crises?: the politics and policy of bank regulation (Princeton University Press, 2008).
23. Haldane, A. G. Rethinking Financial Networks. Speech at "Financial Student Association", Amsterdam (2009).
24. IMF. Global Financial Stability Report: meeting new challenges to stability and building a safer system. Tech. Rep. (2010).
25. Nier, E., Yang, J., Yorulmazer, T. & Alentorn, A. Network Models and Financial Stability. Journal of Economic Dynamics and Control 31, 2033–2060 (2007).
26. Farhi, E. & Tirole, J. Collective Moral Hazard, Maturity Mismatch and Systemic Bailouts. American Economic Review 102 (2012).

27. Eisenberg, L. & Noe, T. H. Systemic Risk in Financial Systems. Management Science 47, 236–249 (2001).
28. Cifuentes, R., Ferrucci, G. & Shin, H. S. Liquidity risk and contagion. Journal of the European Economic Association 3, 556–566 (2005).
29. Rogers, L. C. G. & Veraart, L. A. M. Failure and rescue in an interbank network. Management Science 59, 882–898 (2013).
30. Bo, L. & Capponi, A. Systemic Risk in Interbanking Networks. SIAM Journal on Financial Mathematics 6, 386–424 (2015).
31. Haldane, A. G. The dog and frisbee (2012).
32. Laubenbacher, R. & Pareigis, B. Equivalence relations on finite dynamical systems. Advances in applied mathematics 26, 237–251 (2001).
33. Anand, K. & Bianconi, G. Entropy measures for networks: Toward an information theory of complex topologies. Physical Review E 80, 45102 (2009).

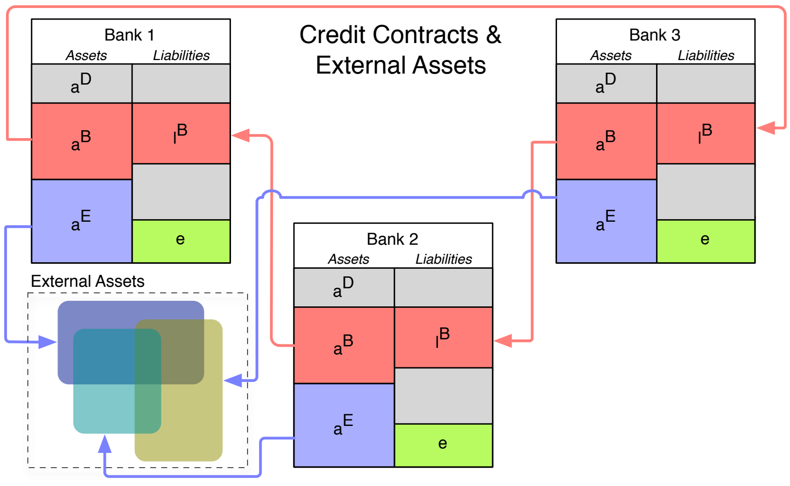

Fig. 1. Example of contract network among three banks. Each bank balance sheet consists of assets and liabilities. The assets of one bank are liabilities of another bank. Credit contracts are represented as arrows (pale red) from the lender to the borrower. Pale blue arrows represent investment in assets issued by entities external to the banking system.

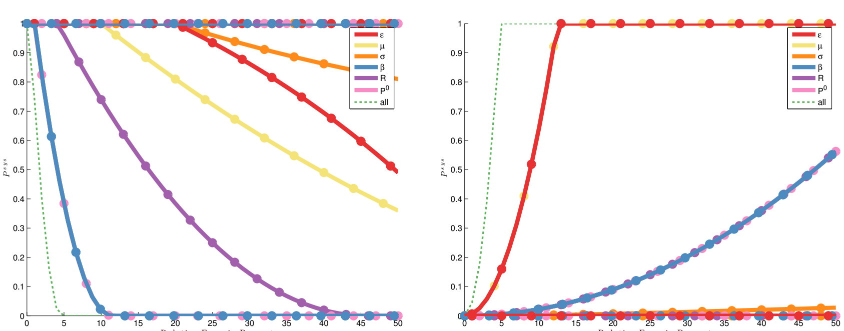

Fig. 2. Systemic default probability vs. relative error on the contract's characteristics. Each pair of curves of a given color represents the minimum and maximum values of the default probability as a function of the relative error on one given parameter (see legend). (Left) For instance, with an error on R (purple curves) larger than 20%, the default probability can take any value between 0.4 and 1. In fact, the maximum value of default probability is 1 for all the parameters when the error is large enough. The green dashed curve refers to the case in which all parameters at the same time contain a given relative error. Shocks are uniformly distributed. Parameter values: $\beta = 3$, $\epsilon = 10$, $\sigma = 0.005$, $R = 0.5$, $P^0 = 0.1$, $\mu = -0.08$. (Right) The maximum probability is one in this case. For instance, with a 10% error on β, the default probability can take any value between 0 and 1. Shocks are uniformly distributed. Parameter values: $\beta = 3$, $\epsilon = 10$, $\sigma = 0.005$, $R = 0.2$, $P^0 = 0.4$, $\mu = -0.01$.

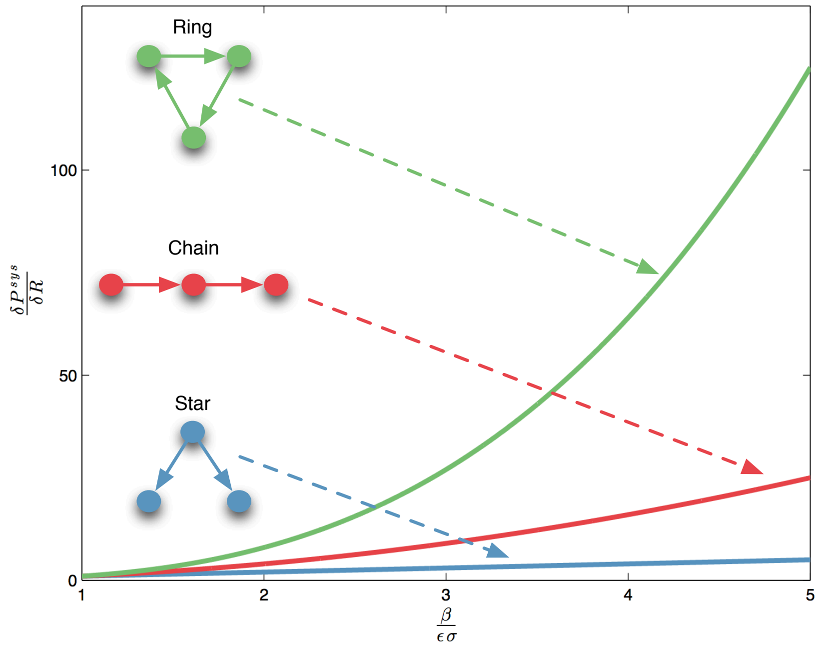

Fig. 3. The sensitivity of the probability of systemic default is amplified by the network structure. Curves represent the sensitivity as a function of the ratio between the interbank leverage and the maximal loss on the external assets.

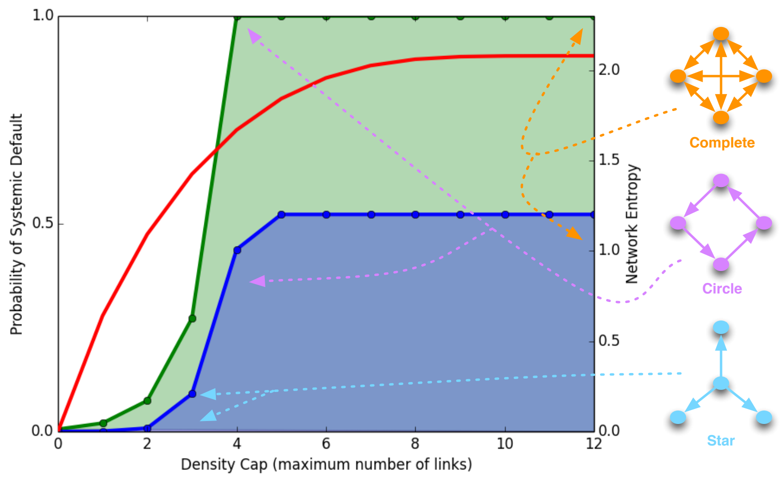

Fig. 4. The probability of systemic default as a function of the maximal network density. We consider all feasible network configurations for a given cap on link density: the green and blue areas represent the range of possible values of probability of systemic default for two parameter sets. Three benchmark configurations are highlighted: out-star, in-star, circle and complete graph. The red curve represents the network entropy. Shocks are uniformly distributed. Parameter values: (blue area) $\beta = 3$, $\epsilon = 10$, $\delta = 3$, $\sigma = 0.08$, $R = 0.5$, $P^0 = 0.1$, $\mu = -0.03$; (green area) $\beta = 3$, $\epsilon = 10$, $\delta = 3$, $\sigma = 0.05$, $R = 0.5$, $P^0 = 0.1$, $\mu = -0.08$.