

The price of complexity in financial networks

Stefano Battiston^{a,1}, Guido Caldarelli^{b,c,d}, Robert M. May^e, Tarik Roukny^{f,g}, and Joseph E. Stiglitz^h

^aFINEXUS - Center for Financial Networks and Sustainability, Department of Banking and Finance, University of Zurich, 8032 Zurich, Switzerland; ^bIMT School for Advanced Studies, 55100 Lucca, Italy; ^cInstitute for Complex Systems, CNR, 00185 Rome, Italy; ^dLondon Institute for Mathematical Sciences, London W1K 2XF, United Kingdom; ^eOxford University, Oxford OX1 1HP, United Kingdom; ^fInstitut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle, École Polytechnique & Solvay Brussels School of Economics and Management, Université Libre de Bruxelles, 1050 Brussels, Belgium; ^gFonds National de la Recherche Scientifique, 1000 Brussels, Belgium; and ^hColumbia University, New York, NY 10027

Edited by Jose A. Scheinkman, Columbia University, New York, NY, and approved July 6, 2016 (received for review November 21, 2015)

Financial institutions form multilayer networks by engaging in contracts with each other and by holding exposures to common assets. As a result, the default probability of one institution depends on the default probability of all of 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 in the probability of systemic defaults. From the point of view of financial regulators, our findings show that the complexity of financial networks may decrease the ability to mitigate systemic risk, and thus it may increase the social cost of financial crises.

financial contagion | financial networks | complexity | default probability | systemic risk

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 because assisting financial institutions in recent years has come with a social cost equivalent to a few percent points of gross domestic product for Organization for Economic Cooperation and Development countries (3). In particular, poor estimates of systemic risk are socially costly because regulators and bank managers end up keeping aside either insufficient or redundant buffers (4–7). One of the difficulties is that financial institutions are connected in multilayer networks both directly, via contracts among each other (loans, bonds, repurchasing agreement, derivatives, etc.), and indirectly, via exposures to common assets (8–15). The default probability of one institution depends therefore on the default probability of all of 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 (16–20) and, hence, requires deeper understanding. In particular, the determination of the probability of systemic events has remained an open problem so far (21–25). Here, we show that, in a network of financial contracts, the probability of systemic default can be very sensitive to errors on information about contracts as well as on information about the complexity of the network structure. 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. Although there may be some intuition for this effect in the literature, it has not been explained analytically or quantified before. Our work aids an understanding of how to contain errors in 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 (OTC), i.e., they are not mediated through a central counterparty. They are also collateralized, i.e., banks have to post collateral 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 Fig. 1 (see *Methods* 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 bank is exposed to, but there can be an error, for instance, in 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 (8). In Fig. 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 in each parameter around a given point. Fig. 2, *Left* refers to the case in which the actual value (i.e., with zero error) of the default probability is 1. 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, Fig. 2, *Right* refers to the case in which the actual value of the default probability is 0. For instance, with an error in 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 in parameters of 10% can lead to an error of 100% in the probability. Note that this result does not imply that the error in the default probability is always large compared with errors in the parameters. However, it shows that there exist cases in which the deviation can be very large. Next, to illustrate how the effects of errors on contract characteristics 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. Fig. 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

Significance

Estimating systemic risk in networks of financial institutions represents, today, a major challenge in both 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.

Author contributions: S.B., G.C., R.M.M., T.R., and J.E.S. designed research; S.B. and T.R. performed research; S.B. and T.R. analyzed data; and S.B., G.C., R.M.M., T.R., and J.E.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. Email: stefano.battiston@uzh.ch.




Fig. 1. Example of a 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.

interbank leverage β and the maximal loss $\epsilon\sigma$ that a bank can withstand on 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*, $\partial P^{sys}/\partial R$ 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 has not been reported so far in

the literature, illustrates concretely the impact of financial complexity on the determination of systemic risk.

Errors in 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 in 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 in 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




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 figure key). (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 of 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: $\beta=3$, $\epsilon=10$, $\sigma=0.005$, $R=0.5$, $P^0=0.1$, and $\mu=-0.08$. (Right) The maximum probability is 1 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: $\beta=3$, $\epsilon=10$, $\sigma=0.005$, $R=0.2$, $P^0=0.4$, and $\mu=-0.01$.

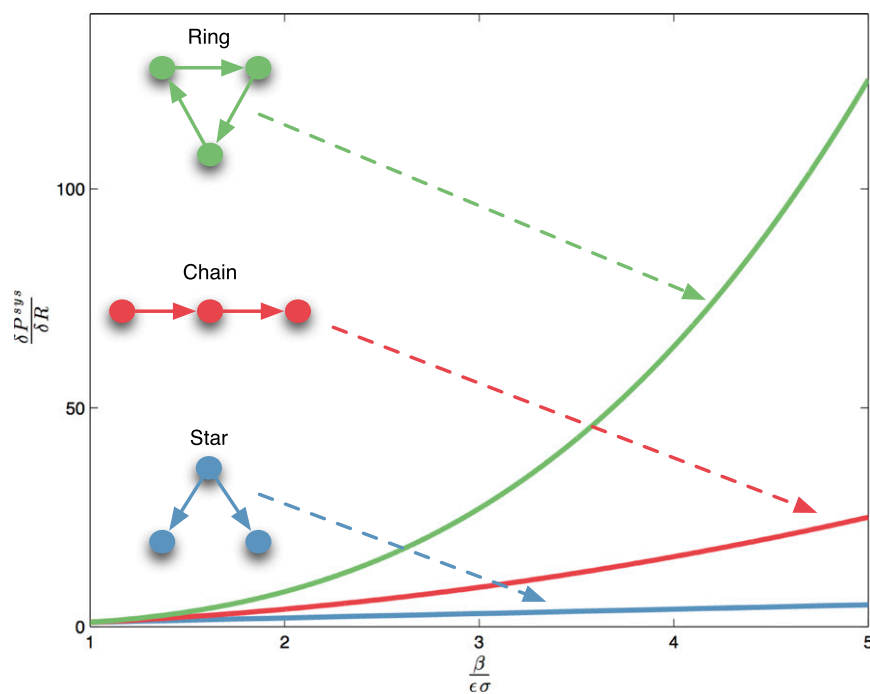


Fig. 3. The sensitivity of the probability of systemic default is amplified by the network structure. Curves represent the asymptotic limit of the sensitivity as a function of the ratio between the interbank leverage and the maximal loss on the external assets.

maximum number of possible contracts, referred to as “density cap” (*Methods*). In Fig. 4, the green area corresponds to a lower expected return on external assets μ , and the blue area corresponds 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 Fig. 4, the red curve represents the entropy $\Sigma(C_i)$ of the space of possible network configurations (*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 (16, 20). 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 bailout 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. However, 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 (8, 26–28) 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 in those works is that the valuation is carried out *ex post* (i.e., at the maturity of the

contracts) and that, in case of default, there is a successful and immediate asset liquidation of the external assets with full recovery. In contrast, asset liquidation implies legal settlements that, in practice, take several months or years, whereas, 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 1 and exogenous (25). Finally, similar to refs. 9, 18, and 19, we intentionally leave aside, at this stage, the question of which configurations and parameters would arise from banks’ dynamic choice, because 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, either a higher complexity of the individual contracts or a higher complexity of the structure of contracts implies that market participants and regulators know less precisely the probability of individual and systemic default. Although 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 the ability to make estimates on default probabilities and hence to correctly price debt instruments. One possible approach to this problem could be to increase the complexity of the regulation in order to match the complexity of the financial market. However, here we have shown that there are intrinsic limitations in the accuracy of estimating the probability of default, which implies that increasing




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, circle, and complete graph. The red curve represents the network entropy. Shocks are uniformly distributed. Parameter values are as follows: (blue area) $\beta=3, \epsilon=10, \delta=3, \sigma=0.08, R=0.5, P^0=0.1$, and $\mu=-0.03$; (green area) $\beta=3, \epsilon=10, \delta=3, \sigma=0.05, R=0.5, P^0=0.1$, and $\mu=-0.08$.

the complexity of regulation may not be an appropriate way to address market complexity (29). Hence, a potential tradeoff emerges between financial stability and market complexity.

Methods

We consider a financial network with 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. Although 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. Fig. 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 , and $x_k^E(2)$ 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 finite variance, the parameter μ_i is the expected return of the portfolio, and σ_i is a scaling factor controlling the effective magnitude of the shocks. We also assume to know the joint probability distribution $p(u_1, \dots, u_n)$.

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 collateralization, 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 ref. 19, although, for the purpose of our study, we exclude rehypothecation, i.e., assets used as collateral are kept aside and cannot be reused. The value of the collateral is assumed to stay constant during the two periods, and it can be ignored in the default condition described in *Default Condition*.

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, and $x_j^B(2)=R_j$ 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 are $a_i^B(2)=a_i^B(1)\sum_j B_{ij}x_j^B(2)$. Notice that this approach differs from previous models (26, 28) that are 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 $e_i(1)>0$, and thus $e_i(2)<0$ iff $[e_i(2)]/[e_i(1)]<0$. It is then convenient to write the default condition as $e_i(1+\mu+\sigma u_i)+\beta_i\sum_j B_{ij}x_j^B(2)-\lambda_i<0$, where the parameter $\beta_i=[a_i^E(1)]/[e_i(1)]$ measures the leverage with respect to external assets, i.e., the magnitude, per unit of initial equity of bank i , of the investments of bank i in external assets (30). Similarly, the parameter $\beta_j=[a_j^B(1)]/[e_j(1)]$ measures the interbank leverage, i.e., the magnitude, per unit of initial equity, of i 's investments in interbank assets (30), and the parameter $\lambda_i=[\ell_i(1)]/[e_i(1)]$ 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's external assets, we finally write the default condition as follows:

$$u_i < \theta_i \equiv \frac{1}{\beta_i \sigma} \left(\lambda_i - 1 - \mu_i - \beta_i \sum_j B_{ij} x_j^B(\chi_j) \right), \tag{1}$$

where θ_i denotes the default threshold. Notice that we have dropped the time in the notation, and 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_j 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

$$\forall i \quad \chi_i = \Theta[u_i - \theta_i(\chi_1, \dots, \chi_n)], \tag{2}$$

where Θ denotes the step function or Heaviside function (i.e., equal to 1 if the argument is positive, and 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' defaults.

The existence and uniqueness of the solution is discussed in *Existence and Uniqueness of Solutions*.

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^p with its expected value $E[x_k] = RP_k + (1 - P_k)$, where P_k is the default probability of bank k . Indeed, expected values are commonly used in 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

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

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

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

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 simplicity, here 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' default conditions are described by a system of nonlinear equations, $\chi = \Theta(u, \chi)$, for which no closed-form solution exists. However, finding all of the solutions of such a system is equivalent to studying, 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: $\Theta: \{0, 1\}^n \rightarrow \{0, 1\}^n$. Because the map is deterministic and the set is finite, for any given initial condition, the map must have either a unique fixed point or a limit cycle; otherwise the set would have to be nonfinite (31). 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 default state of each bank is a nondecreasing function of the default states of the other banks. As a result, 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 $\Delta_p = (\hat{p} - p^*)/p^*$, i.e., the relative deviation of the regulator's estimation, \hat{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 interbank leverage β (i.e., $\Delta_\beta = 0.1$) will

yield $\Delta_\beta \in \{-10\%, -9\%, \dots, 0, \dots, 10\%\}$. For each value of Δ_β , we compute the corresponding probability of systemic default using $\hat{\beta} = \beta^*(1 + \Delta_\beta)$, 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 an example the leverage parameters ε and β . As $\hat{\varepsilon}$ and $\hat{\beta}$ 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 $\hat{\varepsilon}$ and $\hat{\beta}$ given a level of deviation $\Delta_{(\varepsilon, \beta)^*}$.

Analytical Example. 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 $[-1, 1]$. 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 $P^{sys, star} = (1/2^3)(1 + \theta_1^+)(1 + \theta_2^+)(1 + \theta_3^+)$; $P^{sys, chain} = (1/2^3)(1 + \theta_1^+)(1 + \theta_2^+)(1 + \theta_3^-)$; and $P^{sys, ring} = (1/2^3)(1 + \theta_1^+)(1 + \theta_2^+)(1 + \theta_3^+)$. Note that $\theta_i^+ \geq \theta_i^-$ because a bank needs more positive shocks on its external assets to survive when counterparties on the interbank have defaulted. It follows that $P^{sys, ring} \geq P^{sys, chain} \geq P^{sys, star}$. It also follows that the sensitivity of the default probability on the recovery rate R depends on the network structure $\partial P^{sys, ring} / \partial R \propto [\beta / (\varepsilon \sigma)]^3$; $\partial P^{sys, chain} / \partial R \propto [\beta / (\varepsilon \sigma)]^2$; and $\partial P^{sys, star} / \partial R \propto \beta / (\varepsilon \sigma)$. In this example, as long as $\beta / (\varepsilon \sigma) > 1$, 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: $\tilde{N}(C_l) = \sum_{i=0}^{C_l} (n(n-1)/i)$. 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 (32), i.e., the logarithm of the number of configurations normalized by the number of nodes, $\Sigma(C_l) = (1/n) \log \tilde{N}(C_l)$.

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 (<https://github.com/troukny/price-of-complexity>) following the instructions provided.

ACKNOWLEDGMENTS. S.B. and G.C. acknowledge support from the European Commission's Global Systems Science (GSS) initiative and the Future and Emerging Technologies programme (FET) through the Grants SIMPOL 610704, DOLFINS 640772, MULTIPLEX 317532, and SoBigData 654024. S.B. acknowledges the Swiss National Fund Professorship Grant PP00P1-144689. T.R. acknowledges the support of Fonds National de la Recherche Scientifique. J.E.S. and S.B. acknowledge the support of the grant on Financial Stability from the Institute for New Economic Thinking (INET).

- Haldane AG, May RM (2011) Systemic risk in banking ecosystems. *Nature* 469(7330): 351–355.
- Bank of England (2013) *A Framework for Stress Testing the UK Banking System* (Bank of England, London).
- Lambert FJ, Ueda K, Deb P, Gray DF, Grippa P (2014) How big is the implicit subsidy for banks considered too important to fail? *Global Financial Stability Report* (Int Monetary Fund, Washington, DC), Chap 3.
- Rochet JC (2008) *Why Are There So Many Banking Crises? the Politics and Policy of Bank Regulation* (Princeton Univ Press, Princeton).
- Haldane AG (2009) *Rethinking Financial Networks* (Bank of England, London). Available at www.bankofengland.co.uk/archive/Documents/historicpubs/speeches/2009/speech386.pdf. Accessed August 5, 2016.
- International Monetary Fund (2010) *Global Financial Stability Report: Meeting New Challenges to Stability and Building a Safer System* (Int Monetary Fund, Washington, DC).
- Nier E, Yang J, Yorulmaz T, Alentorn A (2007) Network models and financial stability. *J Econ Dyn Control* 31(6):2033–2060.
- Elsinger H, Lehar A, Summer M (2006) Risk assessment for banking systems. *Manage Sci* 52(9):1301–1314.
- Beale N, et al. (2011) Individual versus systemic risk and the Regulator's Dilemma. *Proc Natl Acad Sci USA* 108(31):12647–12652.
- May RM, Arinaminpathy N (2010) Systemic risk: The dynamics of model banking systems. *J Roy Soc Interface* 7(46):823–838.
- Anand K, Gai P, Marsili M (2012) Rollover risk, network structure and systemic financial crises. *J Econ Dyn Control* 36(8):1088–1100.
- Acemoglu D, Ozdaglar A, Tahbaz-Salehi A (2015) Systemic risk and stability in financial networks. *Am Econ Rev* 105(2):564–608.
- Elliott M, Golub B, Jackson MO (2014) Financial networks and contagion. *Am Econ Rev* 104(10):3115–3153.
- Peltonen TA, Scheicher M, Vuillemeij G (2014) The network structure of the CDS market and its determinants. *J Financ Stab* 13:118–133.
- Capponi A, Larsson M (2015) Price contagion through balance sheet linkages. *Rev Asset Pric Stud* 5(2):227–253.
- Battiston S, Caldarelli G, Georg C-P, May R, Stiglitz J (2013) Complex derivatives. *Nat Phys* 9(3):123–125.
- Caballero RJ, Simsek A (2013) Fire sales in a model of complexity. *J Finance* 68(6): 2549–2587.
- Arinaminpathy N, Kapadia S, May RM (2012) Size and complexity in model financial systems. *Proc Natl Acad Sci USA* 109(45):18338–18343.
- Gai P, Haldane A, Kapadia S (2011) Complexity, concentration and contagion. *J Monet Econ* 58(5):453–470.

