the different partial snapshots of financial networks that are available. Not surprisingly, the reconstruction of complex networks from partial information is one of the outstanding problems in the field.

Standard methods (such as maximum-entropy algorithms) have so far proved to be of limited effect in this respect, although they can reveal hierarchical structures[7] in a network. The main problem for these approaches is that the connectivity of a real network is, in general, not reproducible. However, an alternative approach[8] makes use of the fact that some of the properties of these systems are stable, at least in a statistical sense. In this way it is possible to tackle the problem by bringing together complex-network modelling, suitable generalizations of some concepts from statistical physics (equilibrium statistical ensembles, for example) and tools from mathematical statistics (such as the maximum-likelihood method). Using these ingredients, the fundamental statistical features of some important complex networks — real and synthetic — have been reconstructed in considerable detail, and the propagation of distress has been explored using models that have only a limited number of parameters[8].

One of the most recent techniques makes explicit use of the so-called fitness model[9]. This model describes all the situations in which there is, or there is expected to be, a strong correlation between the connectivity (the number of links) and a non-topological feature (fitness) for each node where 'fitness' can be the total capital of an institution in a financial network, and is typically Pareto-distributed in real networks.

Even if only a small portion of a system is known, it becomes possible to reconstruct the statistical properties of the whole in some detail. For instance, this is used in the reconstruction of the World Trade Web (WTW), where the nodes are countries characterized by their GDP (ref. 10), and in financial networks of interbank lending, whose nodes are banks characterized by their total volume of exchanges. In both cases, the most important statistical features of the networks have been determined by knowing the connectivity of less than 10% of the total nodes in the networks.

However, all of these methods can reconstruct only macroscopic or statistical properties. A larger initial set of information is needed to recover the actual microscopic configuration of the system — which would be much more useful to a policymaker attempting to take necessary countermeasures in the face of a crisis. The data are certainly out there, and financial institutions should be encouraged to release them by regulators and by governments

Network by proxy

A different but related approach is to reconstruct the network using a proxy for the information that is missing. This is how the 'DebtRank'[11] was computed for financial institutions during the recent financial crisis. DebtRank is a measure of financial centrality in the banking network, taking into account the impact of the distress of a node across the whole network; reciprocal equity stakes are used as a proxy for the unknown — and possibly uncollectable — information on the network of mutual exposures.

A similar method works for the network of credit default swaps (CDS; the buyer of a CDS is compensated by the seller in the event of a loan default) across financial institutions. In the case of CDS, the problem is particularly acute; despite the crucial role of these products in the stability of markets over the last decade, there is rarely information available on the structure these networks. The interdependencies can be represented by computing the cross-correlation of CDS pairs; even considering only the couples of CDS with enough statistics, it is possible to generate useful insight into the stability of the systems.

Irrespective of the approach used, the importance of network reconstruction in the analysis of financial systems is clear. Recent theoretical advances in network analysis and modelling provide crucial tools that analysts and policymakers will be able to use in the evaluation and control of financial systems. ❐

Guido Caldarelli[1,2,3], Alessandro Chessa[1], Fabio Pammolli[1], Andrea Gabrielli[3] and Michelangelo Puliga[4] are at the IMT Alti Studi Lucca, Piazza S. Ponziano 6, 55100 Lucca, Italy, [2]London Institute for Mathematical Sciences, London W1K 2XF, UK, [3]ISC-CNR, UOS Sapienza, Dipartimento di Fisica, Università Sapienza, Piazzale Moro 2, 00185 Roma, Italy and [4]ETH Zürich, 8092 Zurich, Switzerland.*
**e-mail: guido.caldarelli@imtlucca.it*

References

1. Bonanno, G., Caldarelli, G., Lillo, F. & Mantegna, R. N. *Phys. Rev. E* **68,** 046130 (2003).
2. Garas, A., Argyrakis, P. & Havlin, S. *Eur. Phys. J. B* **63,** 265–271 (2008).
3. Kaushik, R. & Battiston, S. Preprint at http://arXiv.org/abs/1205.0976 (2012).
4. Kullmann, L., Kertész, J. & Kaski, K. *Phys. Rev. E* **66,** 026125 (2002).
5. Caballero, R. J. *J. Econ. Perspect.* **24,** 85–102 (Fall 2010).
6. De Nicolò, G., Favara, G. & Ratnovski, L. *Externalities and Macroprudential Policy* SDN/12/05 (International Monetary Fund, 2012).
7. Clauset, A., Moore, C. & Newman, M. E. J. *Nature* **453,** 98–101 (2008).
8. Musmeci, N., Battiston, S., Caldarelli, G., Puliga, M. & Gabrielli, A. Preprint at http://arXiv.org/abs/1209.6459 (2012).
9. Caldarelli, G., Capocci, A., De Los Rios, P. & Muñoz, M-A. *Phys. Rev. Lett.* **89,** 258702 (2002).
10. Garlaschelli, D. & Loffredo, M-I. D. *Phys. Rev. Lett.* **93,** 188701 (2004).
11. Battiston, S., Puliga, M., Kaushik, R., Tasca, P. & Caldarelli, G. *Sci. Rep.* **2,** 541 (2012).

Acknowledgements

The authors acknowledge support from the European FET Open Project FOC 'Forecasting Financial Crises' (No. 255987) and from the Italian PNR project.

The power to control

Marco Galbiati, Danilo Delpini and Stefano Battiston

Understanding something of the complexity of a financial network is one thing, influencing the behaviour of that system is another. But new tools from network science define a notion of 'controllability' that, coupled with 'centrality', could prove useful to economists and financial regulators.

The financial crisis that erupted in 2008 has made plain the shortcomings of old paradigms in economics and finance, and researchers have turned to other disciplines to seek fresh insight. Network science — thoroughly studied in mathematics and physics for decades — has thus made its way into economics, with financial institutions imagined as the nodes of a network, linked by financial flows, contracts or other interactions. Two main kinds of tool have been provided for the study of financial networks: the first is network statistics, which summarize global properties of a network in one (or few) numbers; the other, centrality measures[1], provides a measure of 'importance' of individual nodes. Both have proved very useful for economists and policymakers, for a couple of reasons.

First, these methods supplied, and confirmed, general insight into financial networks. For example, it is now well understood that the importance of an institution may crucially depend on its position in the financial network, rather than on the characteristics of the institution alone. This idea underlies 'macroprudential regulation', which acknowledges that an exclusive focus on the characteristics of individual banks (such as leverage) will not ensure financial stability.

Second, network statistics and centrality measures can inform actions. For instance, by uncovering the general properties of a network, network statistics may reveal systemic fragilities that can be addressed by regulation of capital buffers or portfolio diversification. Similarly, under the (more or less implicit) assumption that 'central' nodes are good 'targets' to influence a network, centrality measures can pinpoint institutions that deserve the particular attention of regulators. In a network of liquidity flows, the importance of a node can be measured by 'feedback centralities'[1,2], to single out the banks that are most relied on by other banks for liquidity provision.

However, the precise link between the centrality of a node — or the properties of a network — and the possibility of exploiting such nodes or properties to control a system has most often remained implicitly assumed, rather than explicity stated. In the literature, conclusions are typically that 'a network with certain characteristics is exposed to certain fragilities', or 'disruption of certain nodes would most severely affect the network'. But the approach has been contemplative, so to speak, aimed at describing rather than controlling the network.

Only very recently has the issue of how to influence networks has been tackled explicitly in network science[3,4]. Suppose we are given a network of 'influences', determining the evolution of the 'state' of a set of nodes. What is the smallest subset of vertices that we should control, to be able to steer the system in any desired direction? As an example (fictitious, but nevertheless concrete), imagine a banking authority that wishes to ensure each bank under its supervision maintains a level of capital above a certain threshold. Assume this authority has the possibility to influence the rate of growth of each bank's capital and that, at the same time, capital levels are also subject to other effects, stemming from the interaction with the other banks. Finally, either because it would like to be as 'light-handed' as possible, or because it lacks a fully detailed picture of interbank interactions, the authority wishes to interfere with as few banks as possible. Which banks should the authority focus on?

The answer to this question, in the form of procedures or algorithms to identify the 'drivers' of a network, comes from the concept of network controllability[3]. Interestingly, it turns out that network drivers are not necessarily the largest institutions in a system: as in the works of a clock, a small cog can move a larger one and eventually the whole machine. From a practical point of view, a convenient property of these algorithms is that they do not require exact knowledge of the strength of links in the network — and this is particularly welcome in the context of economics, as financial networks are marred by measurement difficulties. This characteristic sets the controllability results apart from several centrality measures, which instead depend on minute details of the network.

We will now consider two examples of how these twin concepts of network controllability and network centrality apply to finance, looking at two types of network that are particularly relevant for the economy: interbank lending networks, in which the links between institutions are financial exposures; and large-value payment systems (LVPS), where the links represent liquidity flows or, in simple terms, payments.

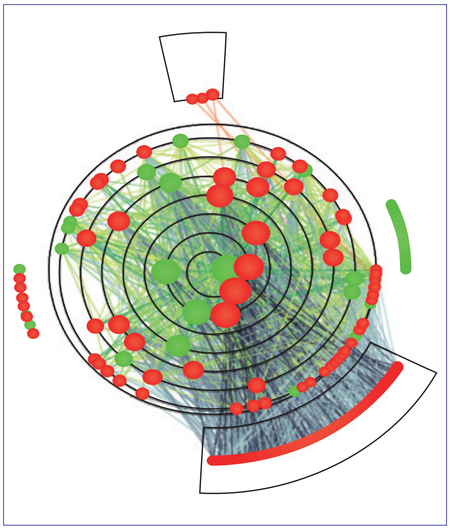

Figure 1 | TARGET2 represented as a 'bow-tie' diagram[9]. This network has 691 nodes (not all represented for the sake of clarity). The higher a node's feedback centrality, the closer it is to the centre of the spiral of the strongly connected component (SCC), and the larger its size. The nodes in the top box represent participants from which liquidity moves out but does not enter, those in the lower right box receive liquidity but do not transfer it back. Some nodes act as drivers (red); others don't (green). Links between nodes within the SCC are colour-scaled from yellow to green; links from a small set of liquidity providers (top) to the SCC are shown in orange; and links from the SCC to a large set of liquidity receivers (lower right) are grey. In all cases, the darker the hue, the higher the degree of the node that the link originates from.

Payment systems

Banks and other large financial institutions settle reciprocal obligations in cash through LVPS. Unknown to most outside the banking profession, and rarely mentioned in the press, LVPS funnel huge liquidity flows across hundreds of institutions every day. The pan-European payment system TARGET2 processes on average 350,000 payments, with a value of €2.3 trillion daily. The corresponding US Fedwire processes 500,000 transactions worth US$2.6 trillion and in the UK, the Clearing House Automated Payment System (CHAPS) sees about 120,000 transactions, with a value of £250 billion — again, daily.

Most LVPS share some key characteristics. The medium accepted in the system is 'central bank liquidity' — that is, deposits held at the central bank of issue, issued by this latter by its legal power to 'create' money. Moreover, these systems work on the 'real-time gross settlement' modality, whereby a payment is legally discharged only when the full amount is transferred from payer to payee. Although the amounts transacted are staggering, as outlined above, the liquidity needed to settle these payments can be much smaller, as banks can 'recycle' liquidity: the liquidity sent by A to B may then be used immediately by B to pay C, and so on. Nevertheless, because the values processed in a day typically exceed by orders of magnitude the liquidity available at any time in the system, it is essential that recycling is carried out in a fast, efficient and, in a sense, cooperative way[5–7]. To avoid liquidity risk and gridlocks, central banks closely oversee LVPS for which they provide the key lubricant: central bank liquidity.

LVPS naturally lend themselves to a network representation. Figure 1 shows TARGET2, globally the largest LVPS by transacted values. The network structure ('pruned'[8] of the smallest links and nodes

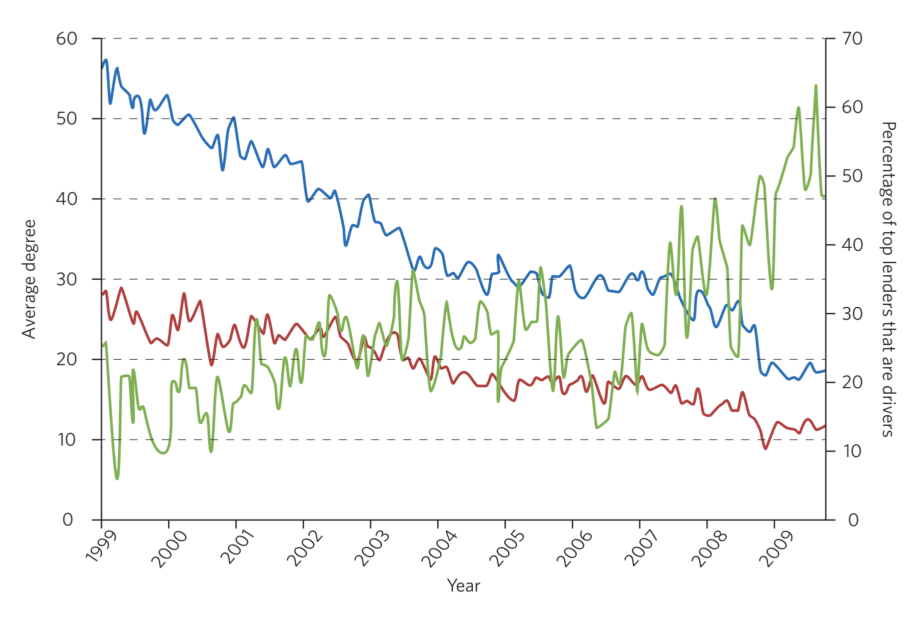

Figure 2 | Transactions of the Italian electronic interbank market e-MID. The average degree of driver banks (red line, left-hand axis) is systematically lower than the average degree of the network (blue line, left-hand axis). Moreover, the percentage of top lender banks that are also drivers (green line, right-hand axis) is well below 100%, but rises significantly at the onset of the global financial crisis in 2008.

for graphical convenience) is highlighted by the so-called bow-tie decomposition[9]. At the top is a set of system participants from which liquidity moves out but does not enter. This group is providing liquidity to a core set of banks — the 'strongly connected component' — which essentially churns large amounts of liquidity within itself. At the lower right, there is a set of participants that receive liquidity but do not transfer it back. The bow-tie diagram illustrates an interesting feature: LVPS may rely on very few banks as liquidity providers.

The concepts of centrality and controllability define the colours and the position of the nodes in the diagram. Centrality (here measured by feedback centrality[1]) determines the radial distance of a bank from the centre of the plot: a 'central' bank appears closer to the centre of the spiral. If a bank can act as a driver for the system — according to the criteria set out for network controllability[3] — then its representative node is coloured red; otherwise, it is green.

Red nodes drive the system's behaviour, in terms of the banks' propensity to make timely payments, and assuming that bank-to-bank influences are proportional to the transacted values.

For TARGET2, it is clear that analysis of centrality and controllability provide complementary information: there is no evident correlation between the centrality of a node and whether or not it is a driver of the system.

Interbank lending

Centrality and controllability analysis can also be usefully applied to networks of credit provision. Just as for LVPS, in a network of interbank loans[10] a bank's willingness to lend money is influenced by its own lenders: when a bank experiences difficulty in accessing funding, it is likely to reduce the amount of credit it will provide to others. It is reasonable to think that driver nodes in the interbank loans network are the most important lenders; surprisingly, we find that this is only partially true.

The network of interbank loans in the Italian electronic market e-MID over the period from 1999 to 2009 is an interesting case in point. As shown in Fig. 2, the network changed over time, and so did its controllability features. However, a constant was that the average degree (that is, the number of connections to other nodes) of the driver nodes was systematically smaller than the average degree of all nodes in the network. Furthermore, not all top lender banks were drivers: in the early 2000s, only 10–25% of top lenders were drivers, a percentage that grew rapidly at the onset of the global financial crisis (when total lending fell abruptly), but remained well below 100%. These results hint at the fact that controllability theory may reveal aspects of a network that are not captured by classic measures of importance.

Getting better

Network science and statistical physics have contributed several tools to describe complex networks in a telling way, and to rank their components according to various notions of 'importance'. These methods are increasingly being taken up by researchers in finance and economics, with a view to understanding financial infrastructure, such as LVPS, assessing the systemic consequences of defaults and contributing to the regulation of financial institutions. For instance, the concepts of being 'too connected' and 'too central' to fail are becoming prevalent in financial regulation, as well as in the scientific debate.

The idea of network controllability is relatively new, and is now finding its first applications in a financial context. Although the theory is probably not yet ripe for application to actual policy design, it is already able to offer fresh insight— as we have discussed. Most interestingly, the use of controllability analysis demonstrates a change in the attitude of economists and scientists alike towards financial networks: the pressing objective is no longer to just observe them, but to influence them. ❐

Marco Galbiati is at the Bank of England, Threadneedle Street, London EC2R 8AH, UK, and formerly an overseer of the TARGET2 system at the European Central Bank, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany. Danilo Delpini is in the Department of Economics and Business (DiSEA), University of Sassari, via Muroni 25, 07100 Sassari, Italy. Stefano Battiston is at ETH Zürich, 8092 Zurich, Switzerland.*
**e-mail: marcogalbiati.london@gmail.com*

References

1. Nicosia, V., Criado, R., Romance, M., Russo, G. & Latora, V. *Sci. Rep.* **2,** 218 (2012).
2. Battiston, S., Puliga, M., Kaushik, R., Tasca, P. & Caldarelli, G. *Sci. Rep.* **2,** 541 (2012).
3. Liu, Y-Y., Slotine, J-J. & Barabási, A-L. *Nature* **473,** 167–73 (2011).
4. Nepusz, T. & Vicsek, T. *Nature Phys.* **8,** 568–573 (2012).
5. Angelini, P. *J. Bank. Financ.* **22,** 1–18 (1998).
6. Bech, M. L. & Garratt, R. *J. Econ. Theory* **109,** 219–198 (2003).
7. Galbiati, M. & Soramäki, K. *J. Econ. Dynam. Contr.* **35,** 859–875 (2011).
8. Glattfelder, J. & Battiston, S. *Phys. Rev. E* **80,** 036104 (2009).
9. Yang, R., Zhuhadar, L. & Nasraoui, O. in *Computational Social Networks: Tools, Perspectives and Applications* (eds Abraham, A. & Hassanien, A-E.) 143–160 (Springer, 2012).
10. Delpini, D. *et al. Sci. Rep.* (in the press).

Acknowledgements

The authors acknowledge support from the European FET Open Project FOC 'Forecasting Financial Crises' (No. 255987) and from the Working Group on Networks and Systemic risk, part of the Institute for New Economic Thinking's Financial Stability Research Program.