



OPEN

Early-warning signals of topological collapse in interbank networks

Tiziano Squartini¹, Iman van Lelyveld² & Diego Garlaschelli¹¹Instituut-Lorentz for Theoretical Physics, Leiden Institute of Physics, University of Leiden, Niels Bohrweg 2, 2333 CA Leiden, The Netherlands, ²De Nederlandsche Bank, PO box 98, 1000 AB Amsterdam, The Netherlands.SUBJECT AREAS:
INFORMATION THEORY
AND COMPUTATION
STATISTICAL PHYSICS
APPLIED PHYSICS
COMPLEX NETWORKSReceived
8 July 2013Accepted
7 November 2013Published
28 November 2013Correspondence and
requests for materials
should be addressed to
D.G. (garlaschelli@
lorentz.leidenuniv.nl)

The financial crisis clearly illustrated the importance of characterizing the level of ‘systemic’ risk associated with an entire credit network, rather than with single institutions. However, the interplay between financial distress and topological changes is still poorly understood. Here we analyze the quarterly interbank exposures among Dutch banks over the period 1998–2008, ending with the crisis. After controlling for the link density, many topological properties display an abrupt change in 2008, providing a clear – but unpredictable – signature of the crisis. By contrast, if the heterogeneity of banks’ connectivity is controlled for, the same properties show a gradual transition to the crisis, starting in 2005 and preceded by an even earlier period during which anomalous debt loops could have led to the underestimation of counter-party risk. These early-warning signals are undetectable if the network is reconstructed from partial bank-specific data, as routinely done. We discuss important implications for bank regulatory policies.

Financial and banking systems are strongly interconnected networks of institutions exposed to both endogenous and exogenous fluctuations^{1,2}. When defaults occur, they cascade throughout the network and can cause the collapse of an entire system, as dramatically witnessed by the recent financial crisis³. As a consequence, the analysis of economic and financial networks as the propagation channel for distress has received a lot of attention^{4–18}. Much effort has been devoted to the search for regularities in the structure of financial networks, i.e. looking for degree heterogeneity, a core-periphery or a modular structure^{8–12}. Similarly, null models have been introduced in order to understand whether part of the observed topological complexity can be explained relatively simply in terms of the observed heterogeneity of vertices^{13–17}. For interbank networks specifically, a lot of attention has been devoted to quantifying the level of *systemic risk* (the risk of the collapse of the system as a whole) determined by a particular network topology, as opposed to the traditional measures of risk defined for individual banks^{18–21}. It turns out that the minimization of (standard measures of) individual risk can often increase the level of systemic risk, which in turn can hurt individual financial entities^{1,2}. This highlights the inadequacy of traditional models and regulation and suggests an analogy with ecological networks^{22,23}.

All the above approaches focus on the structural properties (e.g. systemic risk), associated with a given, static network topology. However, interbank networks are highly dynamic. As distress starts to propagate, existing financial connections might dissolve and new ones might materialize, modifying the way further reverberations of a crisis are channeled through the network. Therefore, one needs to take into account both the (expected) effects of network topology on the stability of the financial system and the reverse effects of (realized) defaults on the structure of interbank networks. This has led to models of interbank networks that dynamically adapt to critical events, with a continuous feedback between topology and dynamics²⁴.

In this paper, rather than introducing a theoretical model, we carry out an empirical characterization of the interplay between realized financial stress and the changes in the observed interbank structure. We address two main questions: *does the topology of an interbank network undergo major structural change as a crisis suddenly manifests itself?* And if so, *are there any topological precursors of this structural change, to be used as early-warning signals of the approaching crisis?* Our results indicate that the answer to both questions is affirmative.

Results

Data. For our analysis we focus on the recent global financial crisis, that manifested itself at the end of 2007 and continued throughout 2008³, and on its build-up phase, which is much more difficult to identify. We selected a dataset reporting 44 quarterly snapshots of the Dutch Interbank Network (DIN in the following), starting from



the first quarter of 1998 and ending with the last quarter of 2008⁹. Each snapshot reports the exposures between Dutch banks at the end of the corresponding quarter, and represents them as binary links directed from the borrower to the lender. Our data include the year when the crisis manifested itself in its strongest form (2008) plus the preceding 10 years, arguably the build-up phase. More details about the data are given in the Supplementary Information (SI). Note that, when studying the propagation of defaults, the binary topology of interbank networks plays the primary role. The magnitude of the connections, while surely important, plays mainly a quantitative role. For instance, the existence and uniqueness of a ‘clearing payment vector’ that clears the obligations of all banks after a default only depends on the interbank topology²¹. Moreover, while a weighted network is of course more informative than its binary projection, recent empirical results^{15–17} have shown that the knowledge of a binary property often conveys more information about a real-world economic network than the knowledge of the corresponding weighted property.

Topological signatures of the crisis. We start by looking for possible topological signatures of the crisis. We find that, at the onset of the crisis the size (numbers of vertices N), the connectedness (number of links L) and the link density (or connectance c) of the network (see fig. 1 right) do not show any significant change in their (roughly stationary) trends (see fig. 1 left).

We can also separately consider the density of the core and of the periphery of the network⁹: the ideal core-periphery model (CP model) assumes that core banks are all bilaterally linked with each other, that periphery banks do not lend to each other, and that core banks both lend to and borrow from at least one periphery bank (discussed further in⁹ and SI). From the right panel of fig. 1, we confirm that the core is much denser than the periphery (as by construction should be). However, the core- and periphery-specific densities, exactly as the overall density, show only a slight jump from the end of 2007 to the beginning of 2008. The size of the change is not significant, as it is of the same order as the fluctuations characterizing the entire 11-years time interval. Taken together, the above results show that the size and density of the network (as well as their core-specific and periphery-specific values) are completely uninformative about the crisis.

However, we are going to show that the picture changes if, after controlling for the size and density themselves, we consider higher-order topological properties (dyadic, triadic, and so on). We first focus on the relative frequency or abundances of the three possible *dyadic motifs* in the observed network, i.e. the number L^{\leftrightarrow} of reciprocated (full) dyads, the number L^{\rightarrow} of non-reciprocated (single) dyads, and the number L^{\leftarrow} of empty dyads (see fig. 2). These numbers are informative only after filtering out size and density effects, or even more complicated topological properties. Therefore, here and in

what follows, we compare each measured quantity X with the expected value $\langle X \rangle$ under a *null model* which has some properties in common with the observed network but is otherwise maximally random. More precisely, we introduce *z-scores* (see Methods section) to quantify the deviation between data and null model. Technically, the method we adopt is an analytical and unbiased one²⁵ based on maximum-entropy ensembles of graphs with constraints²⁶ (see SI for details). We stress that the use of a null model is very different from that of a proper explanatory model: throughout the entire paper, we do not aim at introducing a model that accurately reproduces the data. Rather, the null models we define represent different benchmarks, with various levels of complexity, that discount for the immediate effects of certain topological properties treated as constraints. Comparing the data with the predictions of a null model allows us to determine which observed structural properties are not simply explained by the constraint specifying the null model itself. Indeed, our most informative findings will correspond to a *deviation*, rather than an *agreement*, with null models. It should therefore be clear that null models are by construction in-sample, as it would make no sense to control, in one snapshot of the network, for the effects of a topological property observed in a different snapshot. The inherently in-sample nature of null models is very different from the out-of-sample one of explanatory models, where the fit with one snapshot of the data is used to reproduce different snapshots.

In the left panels of fig. 2 we show the evolution of the *z-scores* over time^{23,27–29} for each of the three dyadic motifs, under a null model that controls for the size and density of the network. The *Directed Random Graph* (DRG), the directed version of the Erdős-Rényi random graph model (see SI) is such null model. Note that the *z-score* is a convenient measure as it immediately shows which motifs diverge significantly from the null model (see Methods section). We find that, while the size and density of the network are relatively stable throughout the entire period, all the dyadic *z-scores* show an abrupt jump in 2008. The crisis period (highlighted in ochre in fig. 2) is characterized by a sudden decrease of the abundance of full and empty dyads, and a sudden increase of the abundance of single dyads. Note that, before the crisis, the number of reciprocated dyads is significantly larger than the expected one, while during the crisis it becomes marginally consistent with the null model (i.e., the random graph).

Similarly, the observed abundances of single and empty dyads become consistent with the null model in 2008. Since the total number of links is more or less stable, the net effect we see is that reciprocal connections suddenly ‘decouple’ and fill previously empty dyads, making single dyads increase and empty dyads decrease with respect to their expected abundance level. So the network seems to suddenly evolve from a fluctuating but roughly stationary configuration (with few single dyads and many full and empty ones) to a ‘crisis’

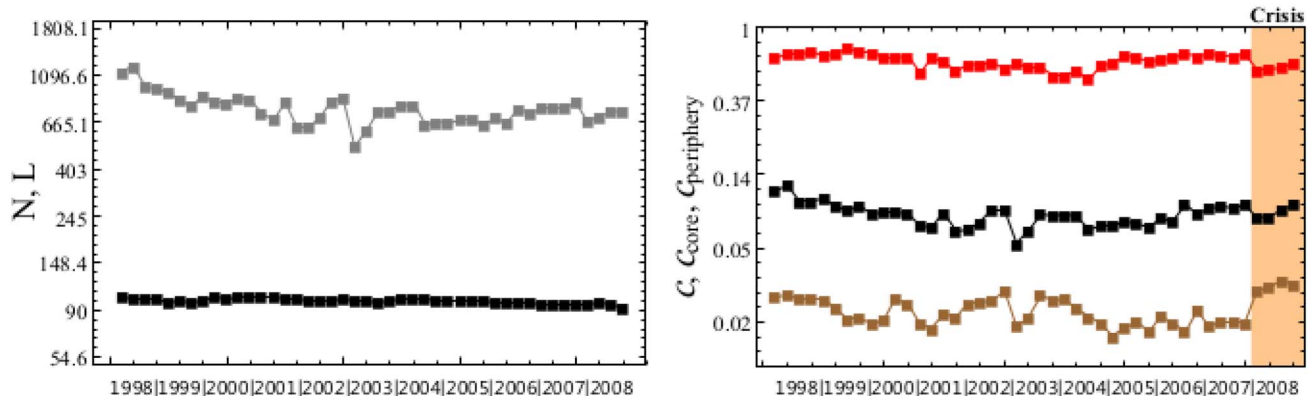


Figure 1 | Left: observed number of vertices (black) and links (gray). Right: observed density of the whole network (black), of the core (red) and of the periphery (brown). The y -scale is logarithmically spaced in both cases.

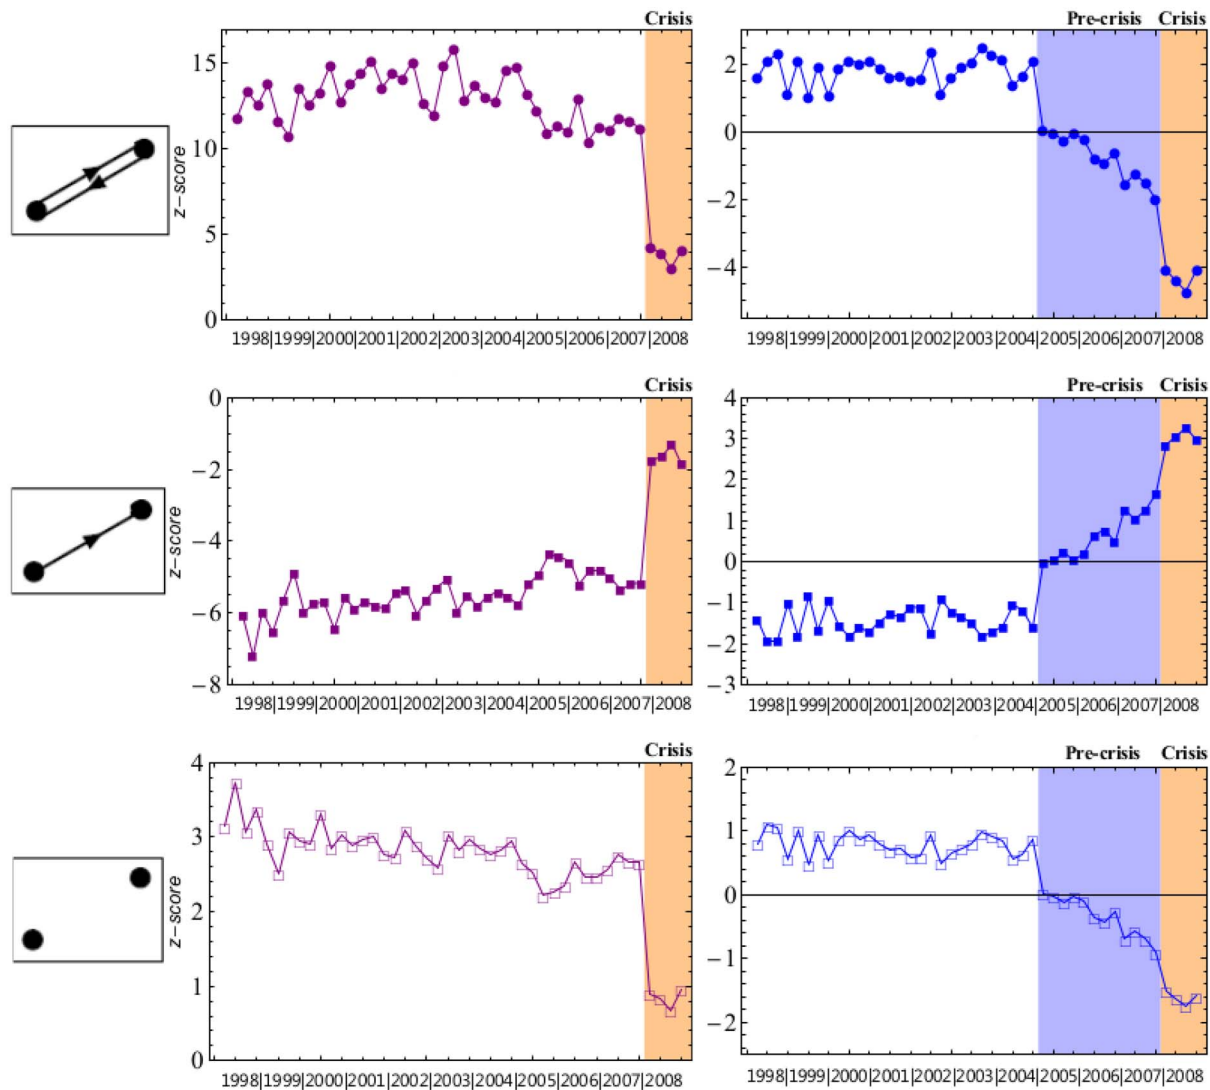


Figure 2 | Temporal evolution of the dyadic z -scores: $z_{L\leftrightarrow}$ under the DRG (top-left, purple circles) and the DCM (top-right, blue circles), $z_{L\rightarrow}$ under the DRG (middle-left, purple, full squares) and the DCM (middle-right, blue, full squares), $z_{L\leftarrow}$ under the DRG (bottom-left, purple, empty squares) and the DCM (bottom-right, blue, empty squares).

configuration whose dyadic structure is marginally consistent with that of an unstructured random graph. We denote this sudden loss of structure as the ‘collapse’ of the original network.

The dyadic motifs, when using the DRG as a reference, are therefore clear topological signatures of the crisis. They are however not predictive, since they show an abrupt transition with no evidence of a preceding build-up phase. They allow us to ‘see’ the crisis, but not to ‘foresee’ it.

Early-warning topological precursors: the pre-crisis phase. Surprisingly, the picture changes completely if we consider a more stringent null model where the intrinsic heterogeneity of banks is accurately controlled for. In particular, we compare each snapshot of the network with a null model (known as the *Directed Configuration Model* (DCM)) where the number of in- and out-going links of each bank is kept equal to the observed values. For a discussion of the properties of the degree distribution(s) of this network, see⁹.

Note that the DRG is an unlikely benchmark economically as the degrees of all banks are narrowly distributed around their empirical average value and banks are thus forced to be similar in size. As a consequence, when studying the deviations of the real network from the DRG, we cannot disentangle the effects of unrealistic bank-specific properties from those of genuine higher-order (dyadic and

beyond) patterns. Incidentally, this shows the main limitation of the *representative agent* concept when applied to economic networks³⁰. By contrast, the DCM indirectly preserves the real heterogeneity of banks, by preserving the observed degrees produced by that heterogeneity. This provides a realistic benchmark with deviations indicating a genuine signature of higher-order effects beyond the bank-specific level, directly arising from the choices of banks.

The second column of fig. 2 shows the dyadic z -scores under the DCM. When comparing these values with the previous ones obtained under the DRG, we find surprising results. Firstly, during the first seven years (quarters 1–28, i.e. from 1998 to 2004 included) all z -scores are stationary and have the same sign as under the DRG, but are much closer to zero. Their small absolute value ($|z| \lesssim 2.5$) suggests that during this period the dyadic structure of the network is not far from the prediction obtained under the DCM, i.e. it is roughly explained by the heterogeneous degrees of banks. By contrast, starting from the 29th quarter (2005Q1), all z -scores suddenly change sign and start to move away from their previous stationary values. This gradually leads to the collapsed network configuration of 2008. The network is then the most distant from the DCM (and, as we observed before, the closest to the DRG). However, the ‘collapse’ is not a sudden structural change, as it is clearly preceded by a 3-year ‘pre-crisis’ period

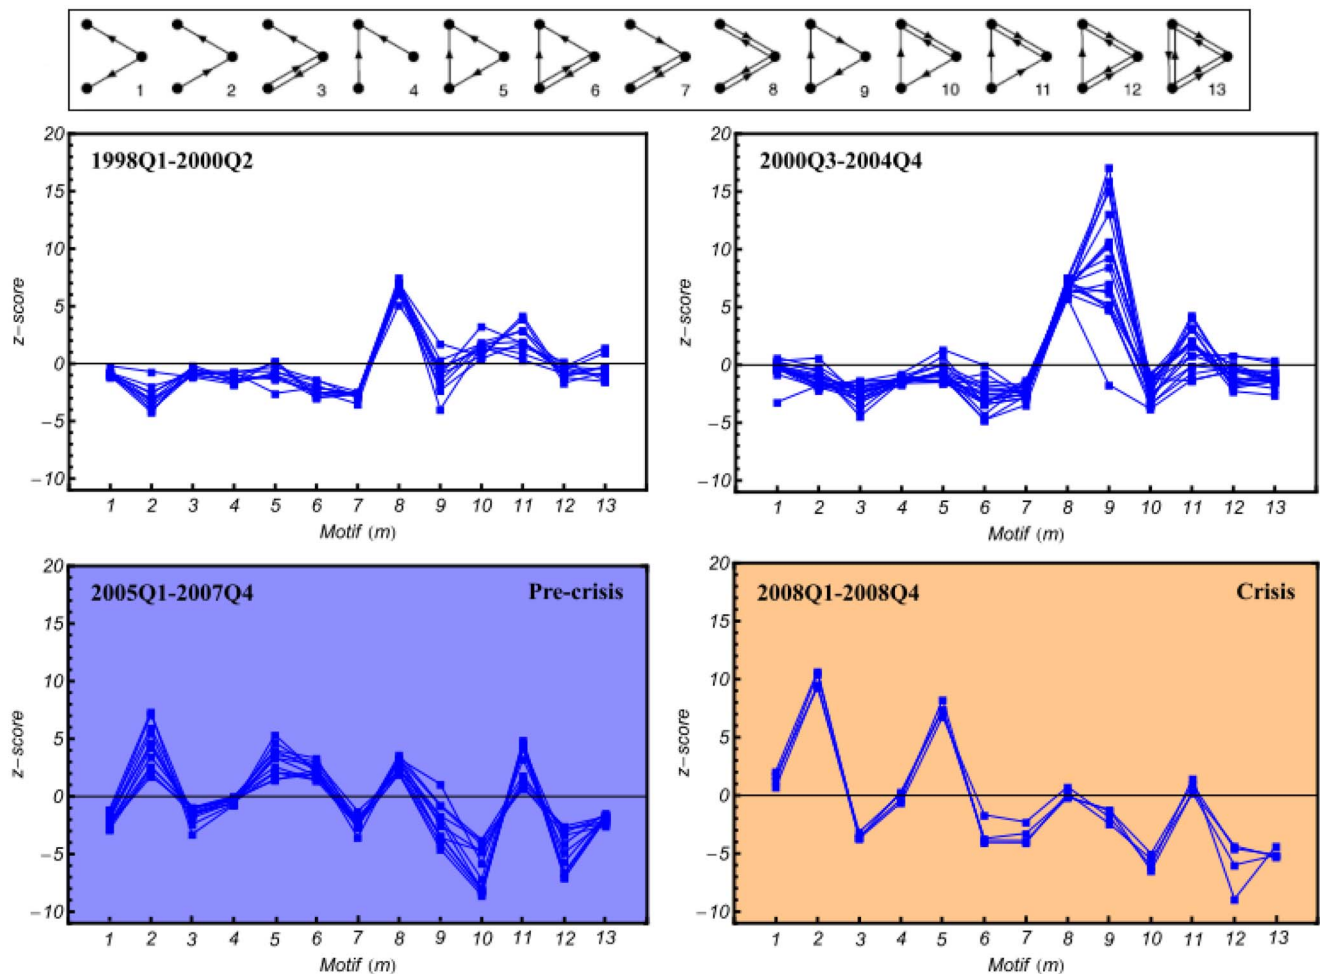


Figure 3 | Triadic z-scores for all the 44 quarters, grouped into four subperiods, under the DCM. First subperiod: t_1-t_{10} (top-left); second subperiod: $t_{11}-t_{28}$ (top-right); third subperiod: $t_{29}-t_{40}$ (bottom-left); fourth subperiod: $t_{41}-t_{44}$ (bottom-right).

(from 2005 to 2007 included, highlighted in purple in fig. 2) bridging the earlier stationary phase to the 2008 crisis.

Our results shown so far suggest that, compared to a homogeneous benchmark, the interbank network displays an abrupt structural transition at the onset of the crisis. In contrast, with a heterogeneous benchmark, the transition is slow and continuous and highlights a gradual build-up phase starting three years in advance of the crisis. The pre-crisis phase is thus an early-warning signal of the upcoming topological collapse.

Topological patterns captured by dyadic motifs are limited to correlations within pairs of vertices. In order to study a higher level of organization, we now analyse the *triadic motifs*^{23,27–29}, i.e. the possible patterns involving three connected vertices.

Exactly as for the dyads, we consider the z-scores for the abundances of each of the 13 triadic motifs (see SI for definitions). However, before considering the development of individual z-scores over time, we first identify the most significant motifs by comparing all 13 z-scores with each other in sub-periods. This results in the ‘motif profiles’^{27,29} shown in fig. 3, where we used the DCM as the null model.

It turns out that the 44 quarterly snapshots do not collapse to a single profile (see SI). By contrast, we can clearly distinguished four subperiods with different characteristic profiles, as evident from the four panels of fig. 3. Remarkably, we find that the last two subperiods coincide exactly with the pre-crisis (2005–2007) and crisis (2008) periods we identified before.

As is clear from the bottom right panel of fig. 3, we can identify the motifs number 2, 5, 10 and 12 as the most significant ($|z| \geq 4.5$) triadic signatures of the 2008 crisis. If we now track these motifs over

time (see left panels of fig. 4), we find exactly the same behavior as shown above for the dyads: the trends over the entire 1998–2004 period are stationary (with small z-scores indicating an approximate accordance with the DCM), and from 2005 onwards they gradually evolve towards the collapsed configuration (for motif 10 the departure actually starts before 2005, but this anomaly will be corrected by a more constrained null model, as we show below). This confirms that the building-up of the crisis is undetectable under homogeneous assumptions, while it becomes manifest in the gradual divergence of the real interbank market from the configuration expected on the basis of the observed heterogeneity of banks.

Clearly, since triads are combinations of dyads, some triadic motifs might be over(under)-represented just because the dyadic motifs they contain are over(under)-represented, in which case the triad as a whole should not be considered an interesting pattern *per se*. In order to control for this, we introduce a more stringent null model that separately controls for the number of single and reciprocated links of each vertex (*Reciprocal Configuration Model* (RCM), see SI for details). The RCM separately preserves the number of empty, full, and single (out- and inward) dyads in which each vertex is involved. As a result, all the observed dyadic abundances are preserved and the dyadic z-scores are zero by construction. In the right panels of fig. 4 we show the triadic z-scores for the same motifs considered previously, but now recalculated under the RCM. We find that motif 2 shows the same trend as before and motif 10, falling in line with it, now confirms the beginning of the pre-crisis phase in 2005. This indicates that motifs 2 and 10 are important building blocks of the network. Motifs 5 and 12 are instead no longer

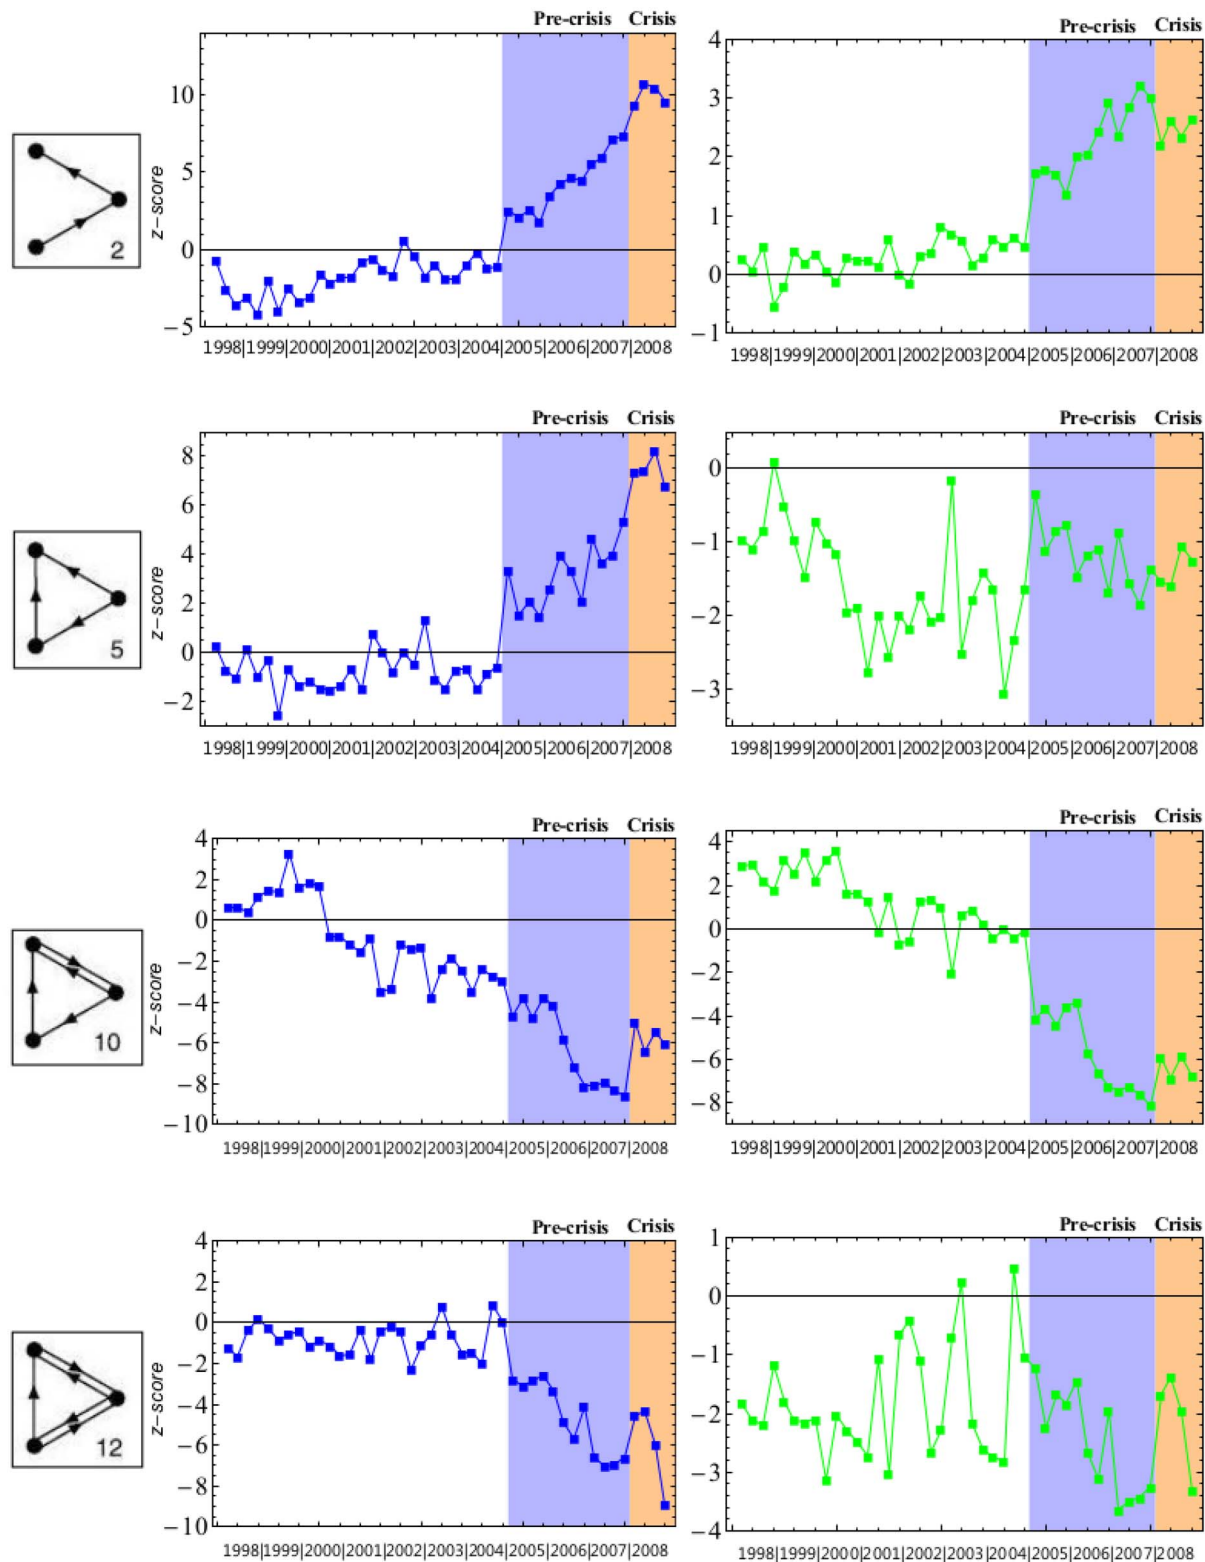


Figure 4 | Temporal evolution of the triadic z-score for motifs 2, 5, 10 and 12 under the DCM (left) and the RCM (right).

significant ($|z| \lesssim 3.5$), and their fluctuating trends do not show any appreciable change during the pre-crisis and crisis periods.

The earliest precursor: anomalous circular lending. What remains to be explained is the nature of the separation (occurring in the mid of 2000) between the first two subperiods shown previously in fig. 3, as all the trends considered so far do not display any significant change at that particular point in time.

Before answering this question, we stress that although fig. 3 might look quite different under the RCM, we find similar results in that case as well. As before, the motif profiles calculated under the RCM over the entire 1998–2008 period do not collapse to a universal distribution (see SI). Still, inside each of the four subperiods we identified earlier, the profiles are coherent (see fig. 5).

Remarkably, the first regime is now almost completely consistent with the null model ($|z| \lesssim 4$ for all 13 motifs), which means that the

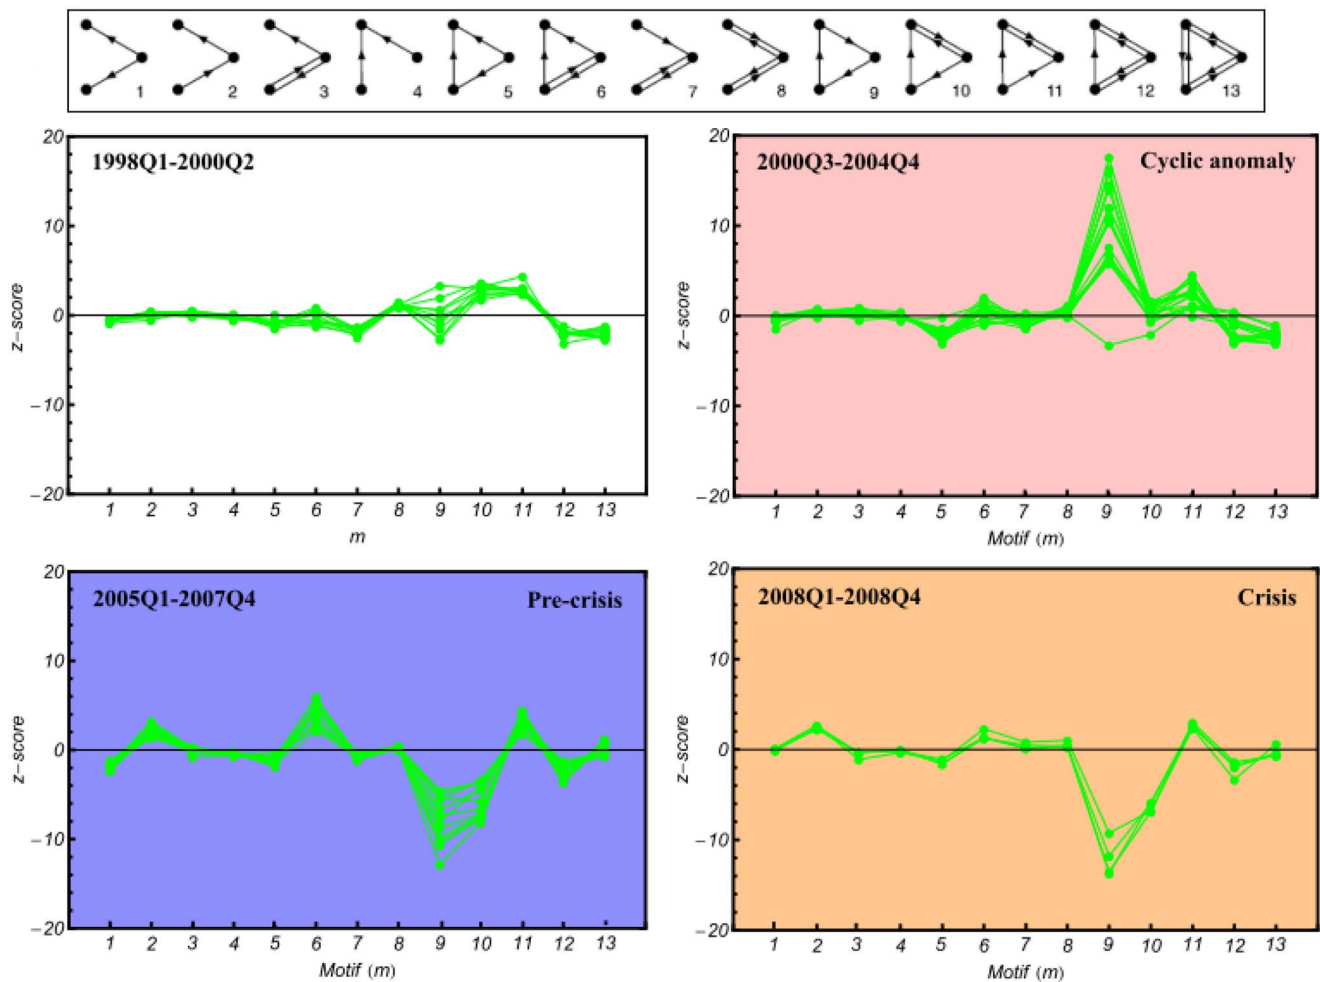


Figure 5 | Triadic z-scores for all the 44 quarters, grouped into four subperiods, under the RCM. First subperiod: t_1 – t_{10} (top-left); second subperiod: t_{11} – t_{28} (top-right); third subperiod: t_{29} – t_{40} (bottom-left); fourth subperiod: t_{41} – t_{44} (bottom-right).

heterogeneous local connectivity and reciprocity of banks entirely explain the triadic structure. Moreover, if we now look more closely at figs. 3 and 5, we find that the main differences between the first two subperiods are determined by motifs 9 (under both the DCM and RCM) and 10 (under the DCM, but not the RCM).

Thus the only significant change occurring in the middle of 2000, after controlling for the dyadic structure, is due to motif 9. The temporal evolution of the latter is reported in fig. 6, under both null models. We find that, from the third quarter of 2000 to the last quarter of 2004, motif 9 indeed shows a marked difference with respect to the rest of the period, and turns out to be strongly over-represented, highlighting an anomalously high number of triads of banks involved in circular lending loops with no reciprocation. Since this subperiod is only characterized by the over-representation of motif 9 (all other motifs are still approximately consistent with the RCM), we denote it as the ‘cyclic anomaly’ phase (highlighted in pink in figs. 5 and 6), and regard it as the earliest precursor of the 2008 crisis. Remarkably, when the cyclic anomaly phase ends and the pre-crisis phase begins, motif 9 suddenly changes from being the strongest over-represented to being the strongest under-represented motif under the RCM (while not significant under the DCM). Thus, it appears that non-reciprocated lending loops, that were the arrangement preferred by triads of banks before 2005, suddenly became the ‘most avoided’ triad. The following two periods (pre-crisis and crisis) are indeed mainly characterized by an increasingly strong under-representation of motifs 9 and 10 (see figs. 4, 5 and 6), which both involve a circular lending loop.

Discussion

The above results have potentially strong implications for bank regulation policies. An immediate one is that the popular view that real interbank markets consist of a well defined core-periphery structure, and consequently that banks can be binarily classified either as big/central or as small/peripheral, is far too simple. Our findings show that the observed heterogeneity of banks is irreducible to the core-periphery dichotomy. Rather, the opposite is true: given the observed heterogeneity of banks, the network is found to have no significant core-periphery structure, and sometimes even has an ‘anti-core’ one (see SI).

The approximate consistency between the real network and the RCM in the initial 1998–2000 period also suggests that, in absence of distress, the topology of real interbank networks might be quite accurately reconstructed using only the knowledge of the number of (inward, outward, and reciprocated) partners of each bank. Technically, this means that, under low stress, real interbank networks might be typical members of an equilibrium statistical ensemble of graphs, where banks’ connectivities are maximally informative. In practical terms, it means that to characterize the network, data requirements are very limited.

However, and more importantly, our findings also show that during the build-up of crises the network can keep moving away from the expectations derived only from the knowledge of bank-specific properties. In this out-of-equilibrium regime, the local connectivities of banks become less and less informative about the network as a whole. This loss of topological predictability speaks against the use of

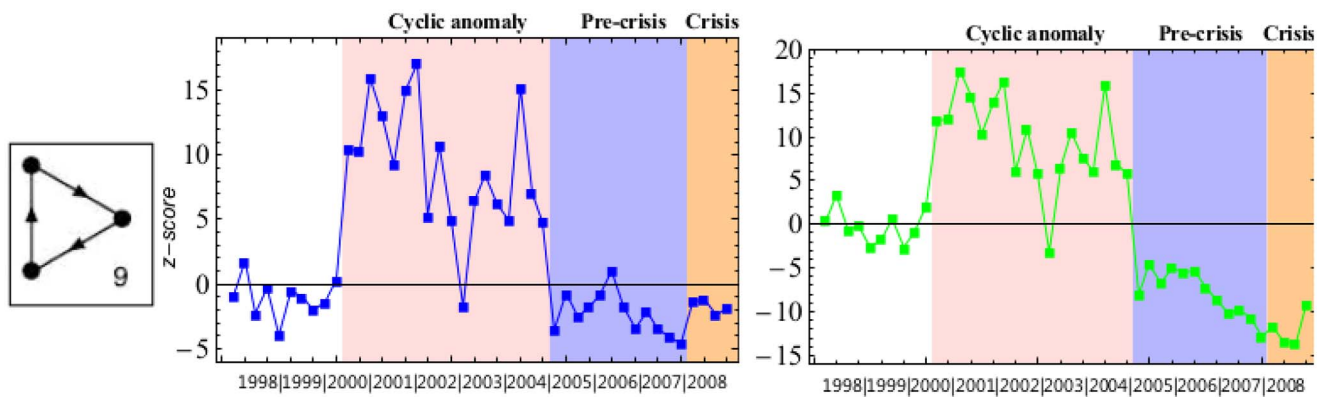


Figure 6 | Temporal evolution of the triadic z-score for motif number 9, under the DCM (left) and under the RCM (right).

maximum-entropy techniques aimed at reconstructing the most likely configuration of an (unobserved) interbank network when only local information about the total assets and liabilities of each bank is available¹⁴. Since assets and liabilities are the (transaction) weighted counterparts of the in- and out-degree of a bank, our results suggest that this technique might yield a realistic guess of the real network only in tranquil times. When the network is under stress, maximum-entropy would instead provide a greatly distorted picture of it.

Strikingly, if our analysis had been carried out on the most likely network consistent with the observed degrees (i.e. the DCM), then every dyadic, triadic, or core-periphery property would have appeared, at each point in time, as perfectly consistent with the configuration model. Supervision based only on bank-specific information, and not on the knowledge of the entire network, is thus likely to remain oblivious to warning signals of structural changes in the run-up to the crisis.

Moreover, both reciprocity and triadic structure have implications for counterparty risk assessment. As other authors have already pointed out^{24,31}, Over-the-Counter (OTC, i.e. not disclosed to third parties) transactions intrinsically generate *risk externalities*: if bank A issues loans to banks B and C, it will require an interest rate that depends on the estimated counterparty risk (which is a function of the fluctuating financial variables on which the ‘health’ of banks B and C depend). But if B issues another loan to C, and A is not aware of this, the interest rate charged by A will underprice counterparty risk, since B becomes vulnerable to a default of C, increasing the correlation between the health of B and that of C. Note that this particular triad is motif number 5. This example shows that the binary topology of interbank networks (more than the intensity of links) has direct effects on systemic risk, and also highlights that some triadic motifs are strongly affected by risk externalities.

Now, it should be noted that the unreciprocated 3-loop (motif number 9) maximizes the underestimation of risk in OTC transactions: each of the three banks involved is not aware of the fact that counterparty risk loops back to itself, creating correlations not incorporated in their bilateral risk pricing. This suggests that, during the cyclic anomaly phase, banks might have systematically underestimated risk externalities.

Note that circularity itself is not necessarily associated with strong risk externalities; but *unreciprocated* circularity is. For instance, within a full dyad, risk loops back between the two banks as well. But in this case both parties are aware of it, and can properly include the increased correlation in their risk pricing. Also note that, while full dyads are still prone to the risk externality involving a third party, this will be a smaller effect since the probability that risk loops back along a longer chain of defaults is smaller than that of risk looping back within the dyad itself. Thus, at a dyadic level, *single* dyads are the most prone to the underestimation of counterparty risk, precisely

because they can become parts of unreciprocated loops. Again, this effect is purely topological: in a mutual dyad with two positive but strongly asymmetric weights, both banks can still incorporate these weights to properly price their risk. Only if one weight is zero, i.e. in a single dyad, this is no longer possible. This further explains why the key information relevant to us is encoded in the binary topology, and not in the intensity of connections. By contrast, at a triadic level, 3-loops involving an increasing number of reciprocated dyads (motifs number 9, 10, 12 and 13 respectively, see SI) are increasingly less prone to the risk externality. Unreciprocated loops (of any length greater than 2) can therefore be considered to be a sort of ‘autocatalytic risk loops’. Since longer loops imply smaller probabilities of cascading defaults, presumably the most dangerous autocatalytic risk loops are precisely those involving three banks.

During the cyclic anomaly phase, all the partly reciprocated loops (motifs 10, 12, 13) were much less abundant than the completely unreciprocated 3-loop (and always consistent with both null models (DCM and RCM)) thus increasing systemic risk. During the pre-crisis phase, the loops with small or no reciprocation (motifs 10 and 9) became increasingly under-represented (figs. 4 and 6). Unfortunately, during the same period reciprocated dyads (that dominated the earlier phase) also became increasingly under-represented, and outnumbered by single dyads (fig. 2). This suggests that, starting from 2005, the underestimation of systemic risk might have progressively increased, first due to autocatalytic risk loops during the cyclic anomaly phase and on a simpler, dyadic level during the pre-crisis phase.

These considerations show that OTC transactions have the potential to create unintentional but emergent, self-reinforcing and destabilizing patterns and feeds into the debate on how OTC markets can be monitored and regulated. Since our results on ‘risk autocatalysis’ suggests that, even when banks spontaneously engage in reciprocated transactions, autocatalytic risk loops can emerge (in the cyclic anomaly phase the reciprocity is still high, cf fig. 5), simply requiring that banks reciprocate a fair amount of transactions (unless this means *all* transactions) is not enough to prevent the creation of unreciprocated loops. One possibility is to introduce a Central Clearing Counterparty (CCP) who would step in the middle of bilateral OTC trades. Although this would reduce the systemic risk due to private interaction, it does introduce the (systemic) risk that the CCP can fail as well. Another approach is to start properly monitoring OTC markets, acting on anomalous motifs. Given the data collection efforts underway for instance in the UK and internationally at the Bank for International Settlements this could be concretely considered (Refs. 32,33).

Although our results are strong on providing early warning signals, they cannot explain the economic rationale for the observed network patterns. As the links are formed in an OTC market, the participants only knowingly create the dyads, not the triadic motifs. A standard explanation is that financial markets are used to hedge risks³⁴. Unexpected idiosyncratic shocks are covered in the market,



either through a cash transaction or through a longer running (derivative) transaction. If things change, which is likely given that shocks arrive continuously, banks generally do not close out a contract but take out a second contract in the opposite direction. Over a reporting period this would lead to significant gross exposure (but much lower net exposure). As during the Great Moderation prior to the crisis, shocks were small, the need to enter into ever more insurance/hedging contracts diminished and thus the probability that any uni-directional link would turn into a reciprocal link became smaller.

However, since the 'phase' transition is taking place both in tranquil and stressed times we need a second explanation for our result: in tranquil times everyone became an acceptable counterparty so hedges no longer needed to be effectuated with existing, known counterparties. Therefore reciprocal dyads slowly became under represented. As the crisis arrived, however, banks seem to actively try to find unconnected parties and thus the existing development towards fewer reciprocal dyads intensified. Such behaviour might be driven by bank's aiming for more diversification, both on the asset and the liability side, achieved by breaking existing links and forming links with hitherto unrelated nodes.

In sum, our results clearly indicate that further theoretical and empirical research is needed to understand the economics of network formation. More generally, any policy directed at regulating interbank markets in a 'pairwise' fashion appears to be fundamentally ineffective, since the most significant patterns are found to occur at an irreducibly triadic level. This result moves the regulation target even further away: while the notion of systemic risk already implies that monitoring individual banks is insufficient to contain systemic risk, monitoring pairs of banks is also likely to fail; the minimal 'building blocks' appear to be triples of banks.

Methods

In order to detect the statistically significant deviations of a measured quantity X from the expected value $\langle X \rangle$, we calculated standard deviation $\sigma[X]$ under the null model and define the z -score

$$z_X \equiv \frac{X - \langle X \rangle}{\sigma[X]} \quad (1)$$

The z -score is a standardized variable measuring the difference between the observed and the expected value in units of standard deviation. If X is normally distributed under the null model, then values within $z = \pm 1$, $z = \pm 2$, $z = \pm 3$ would (approximately) occur with a 68%, 95%, 99% probability respectively. If the observed value of X corresponds to a large positive (negative) value of z_X then the quantity X is over(under)-represented in the data, and not explained by the null model. For most of the topological properties we consider (i.e. dyads and core-periphery structure), the normality under the null model is either trivially ensured by the Central Limit Theorem (CLT), or checked numerically (see SI). The CLT cannot be invoked for triads due to statistical dependencies among the random variables involved (triads necessarily share dyads, and are therefore not independent of each other). Still, larger z -scores identify more significant patterns. Even if null models do not represent, by themselves, a forecasting procedure, it is nevertheless possible to use our analysis to detect a temporal trend, once the deviations highlighted by the z -scores analysis are plotted versus time. The resulting trends, clearly underlining an ongoing structural change, can be interpreted as the starting point of a predictive inference procedure.

- May, R. M., Levin, S. & Sugihara, G. Ecology for bankers. *Nature* **451**, 893–895 (2008).
- Haldane, A. G. & May, R. M. Systemic risk in banking ecosystems. *Nature* **469**, 351–355 (2011).
- Longstaff, F. A. The subprime credit crisis and contagion in financial markets. *J. Financ. Econ.* **97**, 436–450 (2010).
- Schweitzer, F., Fagiolo, G., Sornette, D., Vega-Redondo, F., Vespignani, A. & White, D. R. Economic networks: the new challenges. *Science* **325**, 422–425 (2009).
- Newman, M. E. J. *Networks: An Introduction* (Oxford University Press, Oxford, 2010).
- Goyal, S. *Connections: An Introduction to the Economics of Networks* (Princeton University Press, Princeton, 2009).
- Jackson, M. O. *Social and Economic Networks* (Princeton University Press, Princeton, 2010).
- van Lelyveld, I. & Liedorp, F. R. Interbank contagion in the Dutch banking sector: a sensitivity analysis. *Int. J. Central Banking* **2**, 99–132 (2006).
- van Lelyveld, I. & in't Veld, D. Finding the core: network structure in interbank markets (2012) Dutch National Bank Working Paper. Available at: <http://www.dnb.nl/en/publications/dnb-publications/dnb-working-papers-series/dnb-working-papers/working-papers-2012/dnb276178.jsp> (Accessed: 06/11/2013).

- Liedorp, F. R., Medema, L., Koetter, M., Koning, R. H. & van Lelyveld, I. Peer monitoring or contagion? Interbank market exposure and bank risk (2010) Dutch National Bank Working Paper. Available at: <http://www.dnb.nl/en/publications/dnb-publications/dnb-working-papers-series/dnb-working-papers/dnb233863.jsp> (Accessed: 06/11/2013).
- van den End, W. A. *Credit and liquidity risk of banks in stress conditions. Analyses from a macro perspective* (University Library Groningen, Groningen, 2011).
- Bargigli, L. & Gallegati, M. Finding Communities in Credit Networks, Economics: The Open-Access, Open-Assessment E-Journal **7**, 2013–2017; doi:10.5018/economics-ejournal.ja.2013-17 (2013).
- Bargigli, L. & Gallegati, M. Random digraphs with given expected degree sequences: a model for economic networks. *J. Econ. Behav. Organ.* **78**, 396–411 (2011).
- Wells, S. Financial interlinkages in the United Kingdom's interbank market and the risk of contagion (2004) Bank of England Working Paper. Available at: <http://www.bankofengland.co.uk/research/Documents/workingpapers/2004/WP230.pdf> (Accessed: 06/11/2013).
- Squartini, T., Fagiolo, G. & Garlaschelli, D. Randomizing world trade. I. A binary network analysis. *Phys. Rev. E* **84**, 046117 (2011).
- Squartini, T., Fagiolo, G. & Garlaschelli, D. Randomizing world trade. II. A weighted network analysis. *Phys. Rev. E* **84**, 046118 (2011).
- Squartini, T., Fagiolo, G. & Garlaschelli, D. Null models of economic networks: the case of the world trade web. *J. Econ. Interact. Coord.* **8**, 75–107 (2012).
- Mastromatteo, I., Zarinelli, E. & Marsili, M. Reconstruction of financial networks for robust estimation of systemic risk. *J. Stat. Mech.* **03**, P03011 (2012).
- Upper, C. Simulation results to assess the danger of contagion in interbank markets. *J. Financ. Stability* **7**, 111–125 (2011).
- Battiston, S., Puliga, M., Kaushik, R., Tasca, P. & Caldarelli, G. DebtRank: too central to fail? Financial networks, the FED and systemic risk. *Sci. Rep.* **2** (2012).
- Eisenberg, L. & Noe, T. H. Systemic risk in financial systems. *Manag. Science* **47**, 236–249 (2001).
- Garlaschelli, D., Caldarelli, G. & Pietronero, L. Universal scaling relations in food webs. *Nature* **423**, 165–168 (2003).
- Stouffer, D. B., Camacho, J., Jiang, W. & Nunes Amaral, L. A. Evidence for the existence of a robust pattern of prey selection in food webs. *Proceed. Royal Soc. B: Biol. Sci.* **274**, 1931–1940 (2007).
- Georg, C.-P. The effect of the interbank network structure on contagion and common shocks. *J. Bank. Financ.* **7**, 2216–2228 (2013).
- Squartini, T. & Garlaschelli, D. Analytical maximum-likelihood method to detect patterns in real networks. *New. J. Phys.* **13**, 083001 (2011).
- Park, J. & Newman, M. E. J. The statistical mechanics of networks. *Phys. Rev. E* **70**, 066117 (2004).
- Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D. & Alon, U. Network motifs: simple building blocks of complex networks. *Science* **298**, 824–827 (2002).
- Milo, R., Itzkovitz, S., Kashtan, N., Levitt, R., Shen-Orr, S., Ayzenshtat, I., Sheffer, M. & Alon, U. Superfamilies of evolved and designed networks. *Science* **303**, 1538–1542 (2004).
- Squartini, T. & Garlaschelli, D. Triadic Motifs and Dyadic Self-Organization in the World Trade Network. *Lec. Notes Comp. Sci.* **7166**, 24–35 (2012).
- Kirman, A. P. Whom or what does the representative individual represent? *J. Econ. Perspect.* **6**, 117–136 (1992).
- Acharya, V. V. & Bisin, A. Counterparty risk externality: centralized versus over-the-counter markets (2011) National Bureau of Economic Research Working Paper. Available at: <http://www.nber.org/papers/w17000> (Accessed: 06/11/2013).
- Caruana, J. Interconnectedness and the importance of international data-sharing, Bank of International Settlements Speech. Paper presented at 3rd Swiss National Bank - International Monetary Fund conference on the international monetary system, Zürich (2012/05/08). Available at: <http://www.bis.org/speeches/sp120730.htm> (Accessed: 06/11/2013).
- Langfield, S., Liu, Z. & Ota, T. Mapping the UK interbank system, Bank of England (2012) European Systemic Risk Board and UK Financial Services Authority Working Paper. Available at: http://catedrabde.uji.es/naim/images/presentations/langfield_liu_ota_2012_mapping_the_uk_interbank_system_v110.pdf (Accessed: 06/11/2013).
- Freixas, X., Parigi, B. & Rochet, J. C. Systemic Risk, Interbank Relations and Liquidity Provision by the Central Bank. *J. Money Credit Bank.* **3**, 611–638 (2000).

Acknowledgments

D. G. acknowledges support from the Dutch Econophysics Foundation (Stichting Econophysics, Leiden, the Netherlands) with funds from beneficiaries of Duyfken Trading Knowledge BV, Amsterdam, the Netherlands. This work was also supported by the Netherlands Organization for Scientific Research (NWO/OCW).

The views in this paper are not necessarily those of De Nederlandsche Bank.



Author contributions

T.S. analysed the data and prepared all the figures. I.v.L. provided the data. I.v.L. and D.G. wrote the article. D.G. planned the research. All authors reviewed the manuscript.

Additional information

Supplementary information accompanies this paper at <http://www.nature.com/scientificreports>

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Squartini, T., van Lelyveld, I. & Garlaschelli, D. Early-warning signals of topological collapse in interbank networks. *Sci. Rep.* 3, 3357; DOI:10.1038/srep03357 (2013).



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported license. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/3.0>