



Evaluating structural edge importance in temporal networks

Isobel E. Seabrook^{1,2*} , Paolo Barucca² and Fabio Caccioli^{2,3,4}

*Correspondence:

isobel.seabrook2@fca.org.uk

¹Financial Conduct Authority, Endeavour Square, London, United Kingdom

²Department of Computer Science, University College London, London, United Kingdom

Full list of author information is available at the end of the article

Abstract

To monitor risk in temporal financial networks, we need to understand how individual behaviours affect the global evolution of networks. Here we define a structural importance metric—which we denote as I_e —for the edges of a network. The metric is based on perturbing the adjacency matrix and observing the resultant change in its largest eigenvalues. We then propose a model of network evolution where this metric controls the probabilities of subsequent edge changes. We show using synthetic data how the parameters of the model are related to the capability of predicting whether an edge will change from its value of I_e . We then estimate the model parameters associated with five real financial and social networks, and we study their predictability. These methods have applications in financial regulation whereby it is important to understand how individual changes to financial networks will impact their global behaviour. It also provides fundamental insights into spectral predictability in networks, and it demonstrates how spectral perturbations can be a useful tool in understanding the interplay between micro and macro features of networks.

Keywords: Spectral perturbation; Dynamics; Edge predictability; Classification

1 Introduction

Understanding how individual edges in a network influence its structure and evolution is important in a range of applications. Considering financial networks, network structure has implications for financial stability [1], market efficiency [2] and consumer safety [3]. Identification of players to monitor more closely is of paramount importance to regulators and policy makers, with many attributing the severity of the 2008 crisis to systemic flaws in the banking ecosystem [4].

Our research focuses on understanding how individual edges affect the structure of networks, and how this relates to network stability and evolution. In particular, the purpose of this paper is to show how a spectral perturbation based measure for structural edge importance can be used in the study of temporal networks. We present a brief review of related literature, first considering individual effects on network structure, then those that link network structure to stability and systemic risk, before considering how network structure relates to temporal evolution. We then define a measure for structural edge im-

© The Author(s) 2021. This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

portance l_e ,¹ and we propose a model for network evolution in which an edge's importance can be indicative of future changes. Our results show that l_e values are higher for edges which appear to play a more important structural role, and that subsequent changes occurring in the real networks analysed depend to some extent on the value of l_e .

1.1 Individual effects on network structure

The effect an individual node or edge can have on a network's structure depends not only on the scale of its activity, but also on its position within the network, and the activity of neighbouring nodes and edges. Understanding these interrelations remains one of the key challenges in network science.

Recently, structural node importance has gained a large amount of attention due to its relevance in use cases across a wide range of fields [5]. Methods have predominantly focused on network spectra, in order to illicit structural information from the network adjacency matrix. This includes numerous studies of epidemic processes, in which it is intuitive that the removal of a node that acts as a bridge between communities can be used to stem the spread of a disease, leading to significant effort being taken to understand the influences of community structure on epidemic spreading [5–7]. Similar applications include preventing network-based attacks [8, 9] and understanding and actioning on the spread of gossip in society [10]. This idea of network resilience is often approached from the angle of percolation theory, in which the percolation threshold governing the appearance of a giant component is often related to the leading eigenvalue of the adjacency matrix [11, 12]. An alternative lens is taken by Wang et al. [13] who make use of the observation that the spectrum of the adjacency matrix gives an indication of community structure. In noting that for a network with c strong communities, the c largest eigenvalues of the adjacency matrix are significantly larger than the others, they follow a perturbation based approach to define node importance as the relative change in the c largest eigenvalues of the upon the node's removal. Similar to Wang et al., Lü et al. [14] propose a universal structural consistency index for a network-based on perturbing the adjacency matrix and demonstrate that this index is a good index for link predictability. Restrepo et al. use the same approach to define the dynamical importance of network nodes and edges, instead motivated by the observed relationship between the network leading eigenvalue and dynamical network processes [15]. Our work considers the same central concept of applying edge based perturbations to the adjacency matrix and focusing on the change in the leading eigenvalue, however differs in that we propose the use of this concept as an indicator for subsequent change in networks, as opposed to a measure capturing the effects of node or edge removal on network structure and dynamics.

Many works in the financial literature focus on node specific influence on stability. For example, Battiston et al. [16] define a node ranking coined DebtRank, which takes recursively into account the impact of distress of an initial node across the whole network. Their measure amounts to the fraction of the total economic value in the network that is potentially affected by the distress or default of a specific node. They applied their method to a network of loans from the Federal Reserve to financial institutions between 2008 and 2010, enriched with equity investment relations, and found a strongly connected core of 22 institutions which all became too systemically important to fail at the 2008 crisis peak.

¹Throughout this paper, we use l_e as shorthand for l_{ij} .

They demonstrated the effectiveness of their node ranking in comparison to other centrality measures, and found that it was the only measure to deliver a clear response well before the crisis peak. However, their method specifically considers the case of distress propagation, and does not explicitly measure how an individual node or edge affects the structure of the network in general. Barucca et al. [17] investigate whether a change to few selected banks in the network of the e-MID² market can affect the large scale structure of a network through node removal or degree mutation, and comparing the network structure that results to the original.

Although the bulk of the attention has focused on importance of actors in networks, Helander et al. [18] propose a method for characterising the relative importance of an edge, which they refer to as edge gravity. Edge gravity measures how often an edge occurs in any possible network path. They show that important edges are not necessarily adjacent to nodes of importance as identified by standard centrality metrics, and they also observe that high centrality nodes often have their centrality over-represented by being adjacent to ‘edges to nowhere.’ Similar path-based methods include the BCC_{MOD} (Betweenness Centrality and Clique Model) proposed by [19], which weights the importance of the two nodes forming the endpoints of the edge with the number of cliques containing the edge. Their method outperforms several well-known methods including Jaccard coefficient and betweenness centrality in identifying critical edges both in network connectivity and spreading dynamic. In our work we define the importance of an edge in terms of the change that a small perturbation on the edge would induce in the leading eigenvalue of the weighted adjacency matrix of the network. While other definitions could be considered, we focus here on the leading eigenvalue because it determines for instance the stability of spreading processes on social networks [20, 21], or financial shocks on inter-bank networks [22]. Our methods contrast the above-mentioned path-based approaches by instead considering a network spectrum-based approach, however both approaches show strong connections to node centrality measures; as shown in Sect. 2, an approximation to the network eigenvalue derivative is proportional to the product of the constituent nodes’ centralities. In addition, our research focuses on the temporal behaviour of the network in relation to structural importance, for which future work could consider using alternative measures of structural importance to understand the expected temporal behaviour.

1.2 Network structure in relation to stability and systemic risk

Increasing complexity and stability are inextricably linked, with works as early as May’s investigations into ecosystems with increasing biodiversity highlighting the relationship [23]. In the context of financial markets, although market integration and diversification are widely believed to play a stabilising role [24, 25], Bardoscia et al. [26] demonstrated that two factors of increasing complexity, namely increasing the number of institutions (nodes) and contracts (edges) in an interbank network can drive the system to instability. Similarly, Markose et al. [27] present the idea of institutions being ‘too interconnected to fail’ through an exploration of the structure of the US CDS market. They consider an empirical network constructed from market shares, and make use of the May–Wigner

²e-MID is the Italian electronic market for interbank deposits, a platform for trading unsecured money-market deposits.

condition for stability³ in comparison to a random network. They show that although the CDS structure shows better outcomes than a random network when subject to shocks, the demise of any one big player will bring down other big players. Caccioli et al. [28] showed in a theoretical exploration that uncontrolled proliferation of financial instruments can lead to large instability in markets, and suggest potential interventions such as the introduction of a Tobin tax [29], which is shown by Bianconi et al. to have a stabilising effect [30]. Related to this, Brock et al. [31] used ‘arrow securities’ as a proxy for more complicated hedging instruments, and found that these incentivise construction of larger positions, resulting in a reinforcement effect due to large gains/losses as a result of being on the ‘right’ or ‘wrong’ side of the market. They showed that this is associated with greater instability, and also that the primary bifurcation parameter, marking the onset of instability, occurs earlier when there are more arrow securities. In contrast to the majority of the data centric financial literature which focuses on interbank trading, Bardoscia et al. [22] analysed UK Trade Repository data, which includes all transactions occurring through a Central Counterparty clearing house (CCP) in the UK. Considering a snapshot of the open positions on a single day for interest rate derivatives, FX derivatives and credit default swaps as a three layered network, they compared a ranking derived from the centrality measures to a ranking derived from modelling the network’s response to liquidity contagion, looking at how shocks propagate across the network and translate into payment deficiencies across the different markets. The model considers the stress faced by an institution—the difference between all payments it is required to make and all payment inflows from counterparties, and allows stress to spill over between the layers. They found that centrality measures can be used as a proxy for the vulnerability of financial institutions.

1.3 Network structure in relation to temporal evolution

To understand how networks evolve across time, many researchers have focused on studying the mechanisms for network growth, and defining network models to understand the origin of observed properties of real networks [32–36]. These include the Barabasi–Albert model [37], which demonstrates that scale-free degree distributions observed in real networks can be explained by the presence of growth and preferential attachment in the network evolution. Falkenberg et al. [38] present a simple adaptation to the Barabasi–Albert model, in which new nodes attach to nodes in the existing network in proportion to the number of nodes one or two steps from the target node. This results in an implicit time dependence, which arises from a node’s attractiveness being dependent on its local environment which changes as the network evolves. Central to their model is the idea that network structure and temporal evolution are inherently linked, however their model is limited to the influence of local environment. Others focus on considering temporal networks as multilayer networks, in which one can account for the fact that connectivity patterns in different layers can depend on each other. Bazzi et al. [39] proposed a generative model which explicitly incorporates a user-specified dependency between layers that is flexible enough to incorporate complex interlayer relationships such as dependencies between a layer and all layers that follows, incorporating memory effects into the model.

³The May–Wigner condition for stability is a critical threshold below which any random network has a high probability of stability, and is defined as $D < \frac{1}{Ns^2}$ where D is the network diameter, N is the number of nodes and s is the strength of average interactions between nodes.

A handful of studies have attempted to link global network structure to temporal evolution, such as Peixoto et al. [40], who suggest dynamical variation of the degree-corrected stochastic block model that is capable of finding meaningful large-scale temporal structures in real-world systems and predict their temporal evolution. Their method works with both discrete and continuous time representations, making it versatile to a range of applications. Watts et al. [41] consider semi-random ‘small world’ networks and show that the dynamics are an explicit function of the network structure, and also show find an enhanced propagation speed for small world networks.

A common and general framework for network growth is the fitness model, in which each node has associated with it a time independent ‘fitness’ which represents its propensity to attract links, as proposed by Barabási and Bianconi [42] and further emphasised in [43]. They find that different fitnesses results in multiscaling in the dynamic evolution, or in other words that the time dependence of a node’s connectivity depends on the fitness. Attempts have been made build on this model in order to understand the origins of network dynamics, such as a recent study by Kobayashi et al. [44]. They find that population and activity dynamics are sufficient to explain two types of scaling empirically observed in real networks, however their methods do not explicitly allow for different roles to be captured within a network, by assuming a uniform distribution of fitness parameters. In our research, we explore instead how an edge level quantity derived from the spectrum of a network can similarly be used to determine which edges change in the network. We present methods for estimation of parameters which control both the overall activity in the network, as well as the bias to change for edges with a larger structural importance, and we show how these reproduce behaviours observed empirically.

In the following sections, we look to address two questions: Can we quantify the extent to which an edge affects the overall network structure, and does this provide information on the network’s temporal evolution? We know from the above that network structural information can be gained from the network spectra, both from the observation that the threshold for the appearance of a giant component in a network relates to the leading eigenvalue, and in that the number of communities can be determined from the number of well separated eigenvalues. We also see that the leading eigenvalue provides an indication of stability in terms of dynamical processes occurring on the network. Our aims are to understand edge importance in terms of network structure and stability, so we thus look to capture both of these in our analysis through considering the derivatives of the network’s leading eigenvalue with respect to individual edges. We present evidence that this measure could be a useful indicator in understanding temporal changes in network structure, and we present the results of its application to five real networks. Our main results demonstrate that the elementwise derivative of the leading eigenvalue (l_e) can be predictive of subsequent change for five different networks analysed, and that predictability can be related to the specific realisation of two parameters, α and ρ in the network evolution model in which edges change with probability αl_e^ρ . This has potential implications for stability, as a system experiencing more changes to edges of structural dominance could see a reinforcing effect, leading to an unstable system. These methods could be useful in classifying financial asset systems to inform regulation activities and policy making. We further show that the scale of resultant changes can be related to the realisation of two additional parameters β and γ , again with potential stability implications.

2 Methodology

2.1 Definition of temporal networks

Traditionally, network analytics has focused on static representations of networks, either looking at single snapshots in time, or considering a projection of the time dimension onto a static view by aggregating the links in a time window. In doing so, some, or all, of the temporal information about the network is lost.

However, recently, there have been developments in the modelling of systems as temporal networks, for which the system is represented by a contact sequence (i, j, t) , where i and j constitute the vertex set V at time t . This representation also allows for edges that take time to traverse, or contracts completed after a duration δt by representing the contact sequence as $(i, j, t, \delta t)$ [45]. Since we are considering transactions as instantaneous, we are not interested in transmission time for edges, and we are considering applications where time is discretised, we can formally define a temporal graph $G_t^w(t_{\min}, t_{\max})$ as in [46] as the ordered sequence of graphs, $G_{t_{\min}}, G_{t_{\min}+w}, \dots, G_{t_{\max}}$ where w is the size of the time aggregation (e.g. daily). Element A_{ij}^s of the adjacency matrix at time s is 1 if and only if there exists a link between i and j in G_t , $t \leq s \leq t + w$.

2.2 Central concept—eigenvalue derivatives as a measure of importance

For a given graph $G_t(V, E)$ with adjacency matrix \mathbf{A}^t , the eigenspectrum of \mathbf{A}^t is the set of eigenvalues λ that satisfy the equation

$$\mathbf{A}\mathbf{x} = \lambda\mathbf{x}. \quad (1)$$

By observing changes in the eigenspectrum of a graph, we can gain an insight into structural changes. As we are looking at network snapshots across time, we have a ‘time series’ of graphs and we can consider the change in the leading eigenvalue between successive time snapshots,

$$\Delta\lambda = \lambda(\mathbf{A}^{(t+1)}) - \lambda(\mathbf{A}^{(t)}) \approx \sum_{ij} \frac{\partial\lambda}{\partial A_{ij}} \Delta A_{ij}, \quad (2)$$

where we have made a first order approximation, and the derivative is with respect to the (i, j) th entry of the matrix, as opposed to the entire matrix. Here \mathbf{A} refers to the adjacency matrix, λ refers to the leading eigenvalue of the adjacency matrix and ΔA_{ij} refers to the relative element-wise difference between the two network snapshots, or in other words the change for the individual edge between i and j between the two snapshots.

The two parts of equation (2) can be seen as a payoff between the potential for an edge to influence the structure $(\frac{\partial\lambda}{\partial A_{ij}})$ and the actual change in the network structure (ΔA_{ij}) . Our experiments with synthetic networks look to assess the extent to which our derivation below, which makes approximations and assumptions, captures the true behaviour. The first term measures the sensitivity of the eigenvalue to changes in an individual edge, which we refer to as the structural importance of an edge and denote by l_e . We derive approximations for l_e in equation (3) for the undirected case by taking a perturbation theory approach. Although not explicitly explored in this paper, we also present equation (4) for the directed case. In both cases, we see that the approximations are proportional to the

product of the eigenvector centralities of the two nodes involved in the edge:

$$l_e = \frac{\partial \lambda}{\partial A_{ij}} = 2x_{0,i}x_{0,j}, \quad (3)$$

$$\frac{\partial s^A}{\partial M_{ij}} = \frac{x_{0,i}^M x_{0,j}^M}{2s^A}, \quad (4)$$

where $x_{0,i}$ refers to the i th component of the eigenvector corresponding to the leading eigenvalue, s^A refers to the leading singular value of the adjacency matrix and $x_{0,i}^M$ refers to the i th component of the eigenvector corresponding to the leading eigenvalue of $\mathbf{M} = \mathbf{A}\mathbf{A}^T$. Our definitions are defined in terms of the eigenvector corresponding to the largest eigenvalue, which usually has non-zero values only for the largest connected component of a network. For this reason, in this paper we restrict ourselves to exploring the giant component of the networks, however generalising these to allow for disconnected components will be considered in future work. We note here that our approach is general in that l_e can be computed for all networks, weighted or unweighted, directed or undirected, as differentiability of the spectrum is ensured whenever the adjacency matrix is real and symmetric. The perturbative approach is valid in the case of small, isolated perturbations, which we further explore in Sect. 3.1. Full derivations for these can be found in the Additional file 1, and we validate the approximation for the undirected case in results Sect. 3.1.

We can capture the relationship between l_e and subsequent edge changes by observing the distributions of $P(\Delta A_{ij} = 0 | \ln(l_e))$ and the joint probability $P(\Delta A_{ij}, l_e)$, which we explore in detail in the results Sects. 3.2.5 and 3.3. Our findings from these are compared to our model for the temporal evolution of networks, which we propose in Sect. 2.3, to assess the extent to which our model captures the true behaviour observed.

The second term considers the changes that subsequently occur in response to the value of l_e . This is of significance from a stability perspective; edges that are structurally important could cause a system to become unstable by changing frequently or by a large amount. Conversely, they may also act to stabilise a system if it begins to move towards a regime of instability. This can be explored by assuming that the evolution of our temporal graph is Markovian. We consider this first of all in the proposal of a model for network evolution, parameterised by the extent to which l_e is indicative of the propensity of an edge to change, and the scale of the resultant changes. We further assess the predictability of changes from the value of l_e through the use of a logistic regression classifier, and relate the performance of this to the model parameters.

2.3 Model for network evolution

In order to understand the relation between structural importance and stability of a network over time, we need a model that captures two behaviours. The first of these is that the value of l_e is indicative of the probability for an edge to change, and the second is that the size of a resultant change can be related to l_e .

We thus propose a model in which we can control the extent to which l_e influences a subsequent edge change, both in probability of occurrence and resultant scale. Specifically, we propose a model in which the network evolution exhibits the Markovian property as

in [33, 40]:

$$A_{ij}^{t+1} = \mathcal{V}_{ij}^t A_{ij}^t \mathcal{U}_{ij}^t + (1 - \mathcal{V}_{ij}^t) A_{ij}^t, \quad (5)$$

where \mathcal{V}_{ij}^t follows a Bernoulli distribution $\mathcal{B}(\alpha(l_e)^\rho)$, and \mathcal{U}_{ij}^t is the distribution of edge changes. Here we introduce two parameters which control the probability of an edge to change— ρ which controls the level to which the value of l_e influences the probability for an edge to change, and α scales \mathcal{V}_{ij}^t to ensure that it is a valid probability. A positive value for ρ indicates that more important edges are more likely to change, and a negative ρ would indicate the opposite.

The simplicity of this model means that we are unable to account for edges appearing and disappearing in the network. We will look to incorporate this in future research.

2.3.1 Parameter estimation in real networks

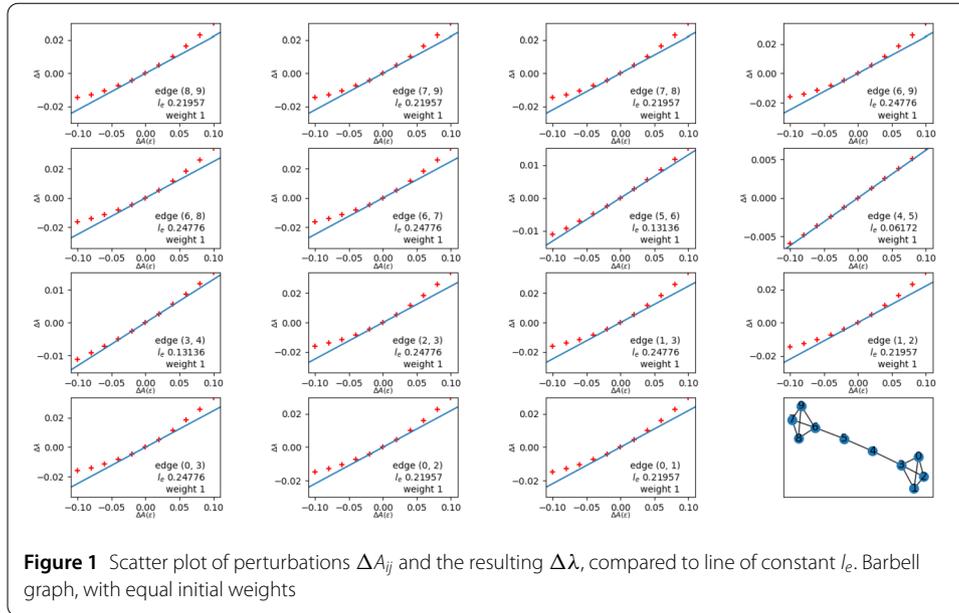
Assuming that our data evolves according to the model in equation (5), we can use observations from real networks to estimate the most likely values of α and ρ from the data. Following a maximum likelihood approach, we can derive estimations for these parameters, by maximising the following log-likelihood as proposed in the Additional file 1:

$$\ln(L(\mathbf{k}|\boldsymbol{\theta})) = \sum_e^N k_e \ln(\theta_e) + (1 - k_e) \ln(1 - \theta_e), \quad (6)$$

where $\theta_e = \alpha l_e^\rho$, and k_e is the observed outcome of edge e . We note here that since α and ρ are constrained to result in a valid probability calculated from αl_e^ρ , the minimisation is subject to constraints and must satisfy the Karush–Kuhn–Tucker conditions [47]. In practice, numerical optimisation of the log-likelihood in equation (6) was used to estimate α and ρ .

2.3.2 Structural influence and network predictability

Depending on the values of the parameters for a given dataset, we might expect the observed values of l_e to be predictive of subsequent change. Specifically, since ρ controls the relationship between l_e and the propensity for an edge to change, a high value of ρ would suggest that l_e would be more predictive of future change. Similarly for α , within the constraints for αl_e^ρ to give the probability of an edge to change, a larger α factor will increase the distance between change probabilities for edges with different l_e , thus also strengthening the relationship between the value of l_e and the propensity for an edge to change. In order to evaluate these effects, we make use of logistic regression for classification of edges into changing vs. unchanging from the values of l_e , and compare the results to a null model consisting of the average over multiple trials in which edges randomly change with probability equal to the fraction of observed changes. The data is split into training and test sets in a stratified manner, with 20% used to test the model on unseen data. The predictions are compared according to balanced accuracy, defined as the average of recall obtained on each class, and Area Under Curve scores for both Receiver Operating Characteristic curves and Precision Recall curves.



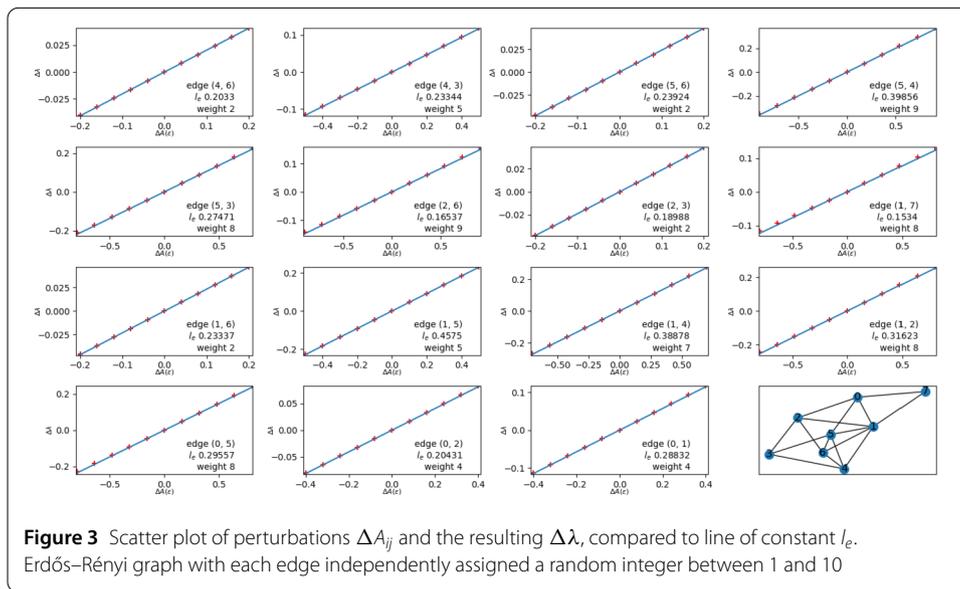
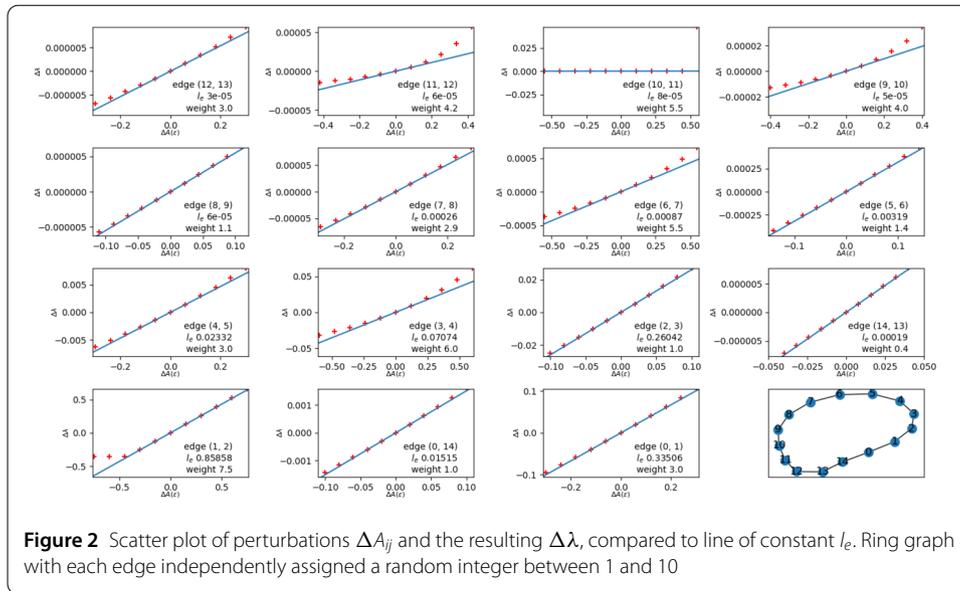
3 Results

3.1 Validation of l_e using toy networks

Here we assess the extent to which the approximations made in calculating l_e hold. We do this by approximating the change in eigenvalues as the coefficient weighted sum of the edge weight changes, $\Delta \lambda = \sum_{edges} l_e \Delta A_{ij}$, and comparing the gradient of this to the value of l_e . Our derivation of l_e makes the simplification in assuming that edge changes occur independently of each other. Our first test thus considers the case of an individual edge changing at each timestep, and we consider perturbations applied to a barbell graph, to observe the effects of network structure, a ring graph, to observe the effects of weight with structural equivalence, and a Erdős–Rényi (ER) graph as a baseline. The results in Figs. 1, 2 and 3 show the line of constant l_e , overlaid with the observed ΔA_{ij} and corresponding $\Delta \lambda$ values.

We see here that our linear approximation generally holds for relative edge changes less than $\Delta A_{ij} = 0.05$. We also see for the barbell graph that l_e captures the structural role of the edges, with edges in the cliques having higher values of l_e than those in the bridge. For the ring graph, we observe a poorer fit for edges with low values of l_e , and the larger l_e edges tend to be adjacent to edges with similar l_e values. Although the edge with the largest weight also has the largest value of l_e , in general there does not appear to be a simple relationship between edge weight, or weight of neighbouring edges, and the value of l_e . For the weighted random network we see similar observations are made for the weighted ER graph, with the lowest l_e values observed for more peripheral edges, and the two edges with the largest weights also having the highest l_e values. Further results for the case of a weighted barbell, and unweighted ring and random networks are shown in the Additional file 1.

Results for the case of two edges changing are also shown in the Additional file 1. In these we observe for the barbell graph better fit is observed for higher values of l_e . For the ring networks and random networks, we see that our model performs well if the observed edge has a larger value of l_e than the other changing edge, but performs poorly when the

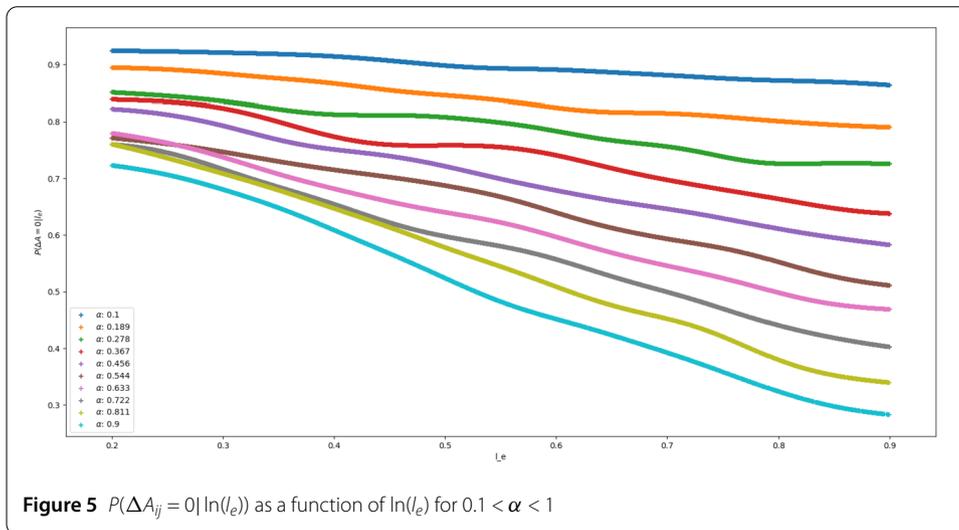
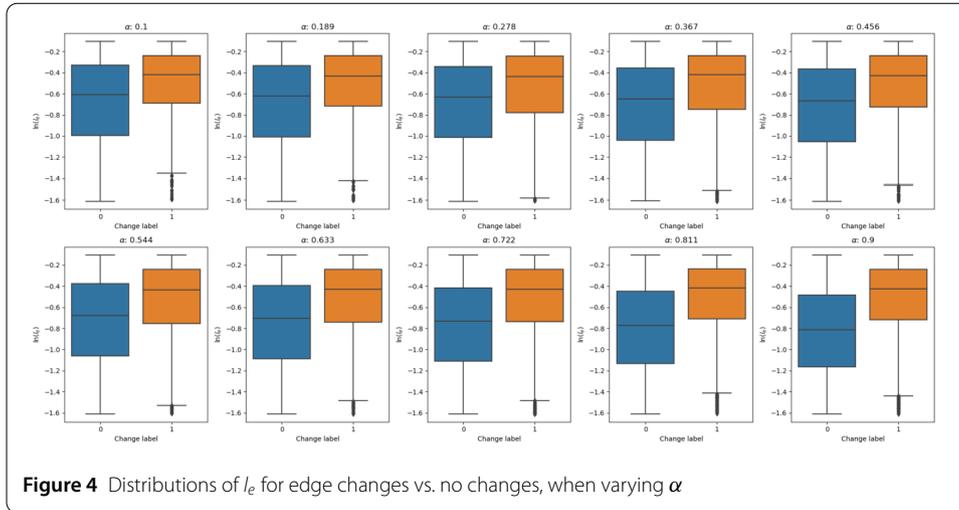


value of l_e is smaller. The case of complete structural equivalence and equal weights in the unweighted ring network shows good performance for all edges.

The breakdown of the method when there are multiple changes occurring between snapshots suggests that our approximation for l_e may be better suited to a continuous or pseudo-continuous representation of a temporal network, which can be seen as the limit of a discrete temporal network in which each snapshot captures an individual edge change occurring at an infinitesimally different time to the neighbouring snapshot changes.

3.2 Relationship between l_e and the presence of edge changes

We can understand the role of the parameters α and ρ by observing the effect of varying the parameters on the distributions of the values of l_e for changing vs. non-changing edges,



$P(\Delta A_{ij} = 0 | \ln(l_e))$. We first consider this for data generated according to our model in equation (5), first keeping ρ fixed and varying α , then fixing α and varying ρ .

3.2.1 Model with varying α

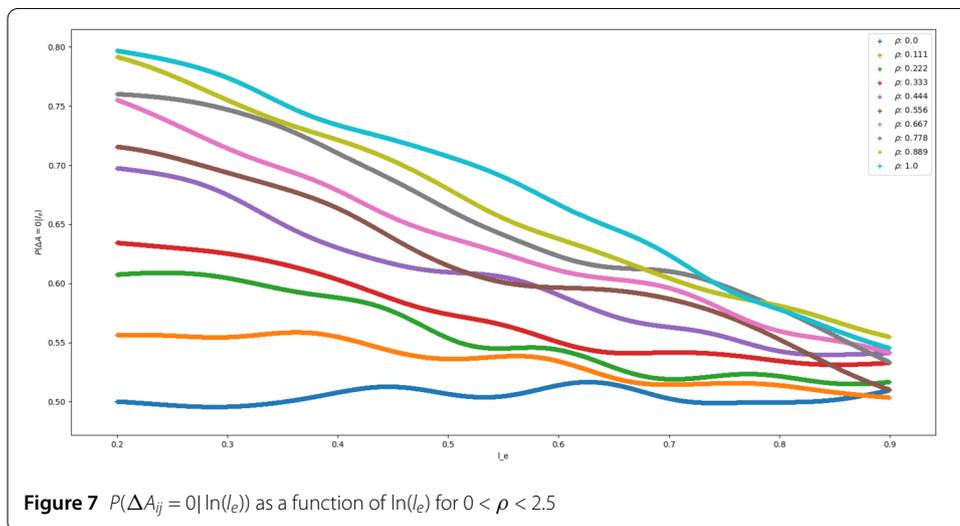
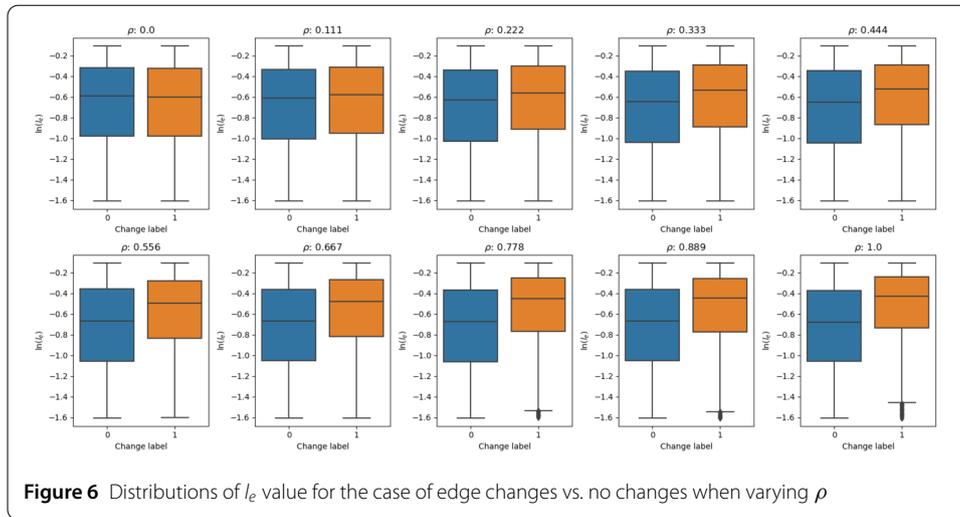
Figures 4 and 5 show the resulting distributions for varying values of α . We see that an increase in α results in a decrease in the probability of an edge to remain unchanged for all values of l_e , and for larger values of α , the rate of increase of change probability with l_e is slightly larger.

3.2.2 Model with varying ρ

Figures 6 and 7 show the distributions when varying ρ . We see here that for increasing ρ , the probability of observing no change increases, and also for increasing l_e , the probability decreases for a given ρ , at a rate that shows a significant dependence on ρ .

3.2.3 Predictability improvement with α and ρ

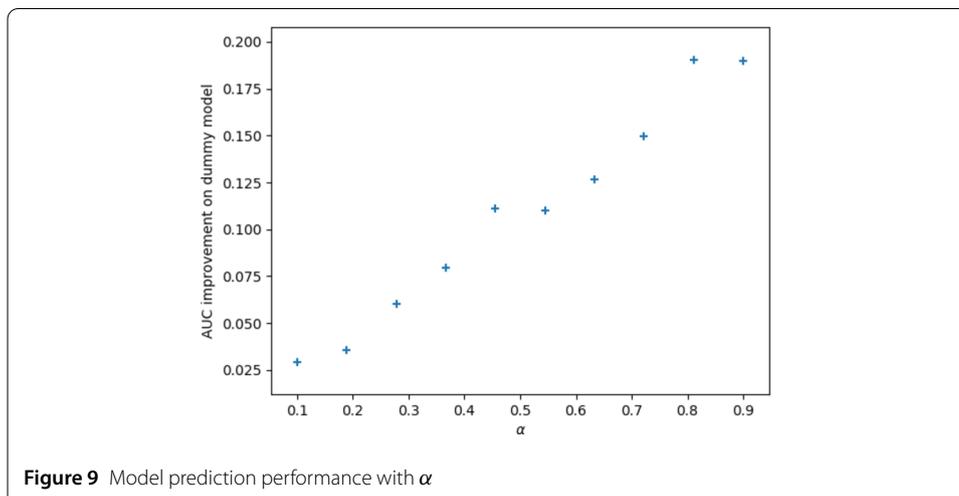
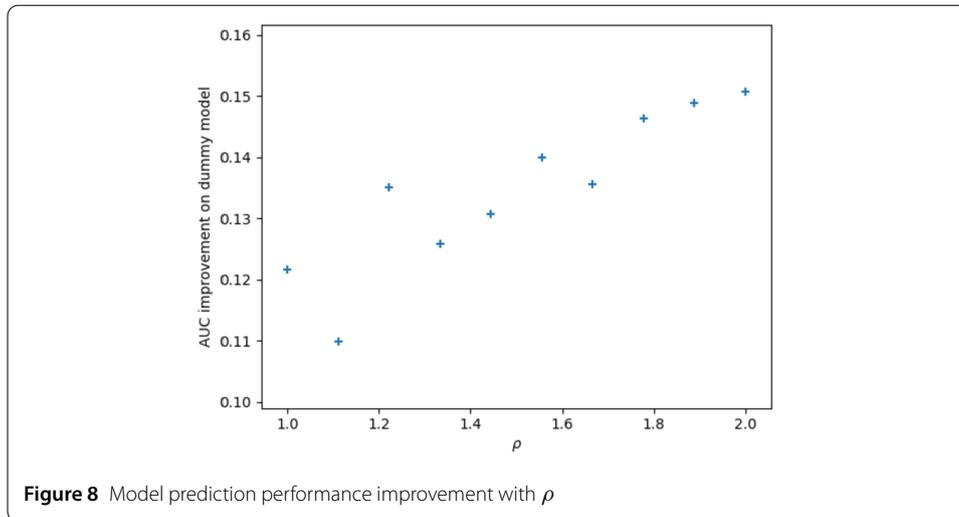
As detailed in Sect. 2.3.2, here we apply a logistic regression classifier with single feature l_e , to datasets with varying α and ρ . Figures 8 and 9 show the improvement in the test



set Precision–Recall Area Under Curve scores for increasing values of each parameter. We see from these that increasing both parameters improves the predictability of changes given the value of l_e , consistent with our observations of the rate of increase of change probability being positively correlated with both α and ρ .

3.2.4 Static observations in real data

We have seen in the above in application to synthetic networks that our model behaves as expected, with networks with a large ρ (and α) being more predictable. Now we explore the performance of our structural influence metric and model through the application to five real datasets. Firstly, given that our research has been motivated by a need to monitor risks in a financial setting, we considered a network of country level bilateral trade [48] and three different capital markets transaction datasets reported under MIFID II regulations. However, our methods can be applied more generally to any temporal networks, and due to the availability and high volume of research conducted into social networks (see [10]), we also considered a network of messages sent between College students [49]. A full description of these can be found in the Additional file 1.



In order to understand the usefulness of l_e as a metric for structural importance, we first examine the edges that rank the highest according to their values of l_e for the bilateral trade dataset, since the historical context of international trade can give us an idea of which edges we might expect to be ‘important’. For the bilateral trade dataset, we see the largest values of l_e for the edge between Portugal and Spain in 1872, and considering the sum across all time, for Greece and Turkey. These are examples of edges with both nodes having large eigenvector centrality; edges involving only one central node are seen to have lower values of l_e . This means that inter-European edges almost exclusively make up the top 100 ranked edges, whereas the lowest ranked l_e edges occur when one, or both, of the nodes have very low centrality scores. Similarly, for the other datasets, the highest values of l_e were also observed for edges involving nodes with high eigenvector centrality. In general, we see that the rankings of l_e are uncorrelated with the rankings of edges according to their betweenness centrality, or their mean value of ΔA_{ij} , however do for some cases correlate with the product of the participating node’s degrees and strengths, as shown in Table 1.

As these datasets contain large numbers of edges (the smallest contained 2785 edges), we cannot fully explore all of the individual observed values of l_e as for the toy networks.

Table 1 Spearman’s rank correlations for l_e with the rank by edge weight, edge betweenness centrality and product of nodes’ degrees

Dataset	$\text{Corr}(l_e, \Delta A)$	$\text{Corr}(l_e, \text{EBC})$	$\text{Corr}(l_e, \text{deg}_{n1,n2})$	$\text{Corr}(l_e, S_{n1,n2})$
Bilateral Trade	-0.061	-0.397	0.352	0.786
College Message	-0.169	0.035	0.581	0.434
Equity-1	-0.104	0.135	0.717	0.320
Equity-2	-0.047	0.166	0.580	0.265
Equity-3	-0.010	0.041	0.923	0.763

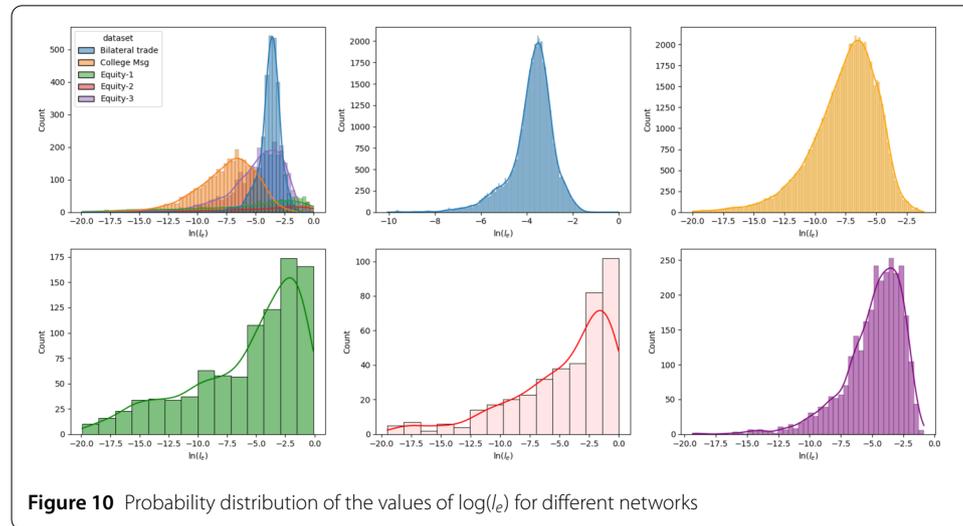


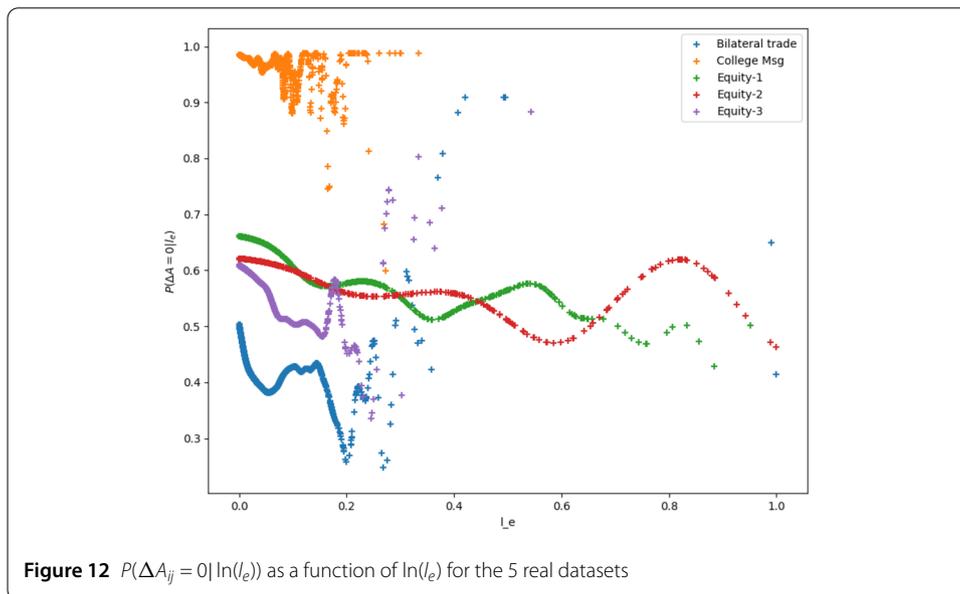
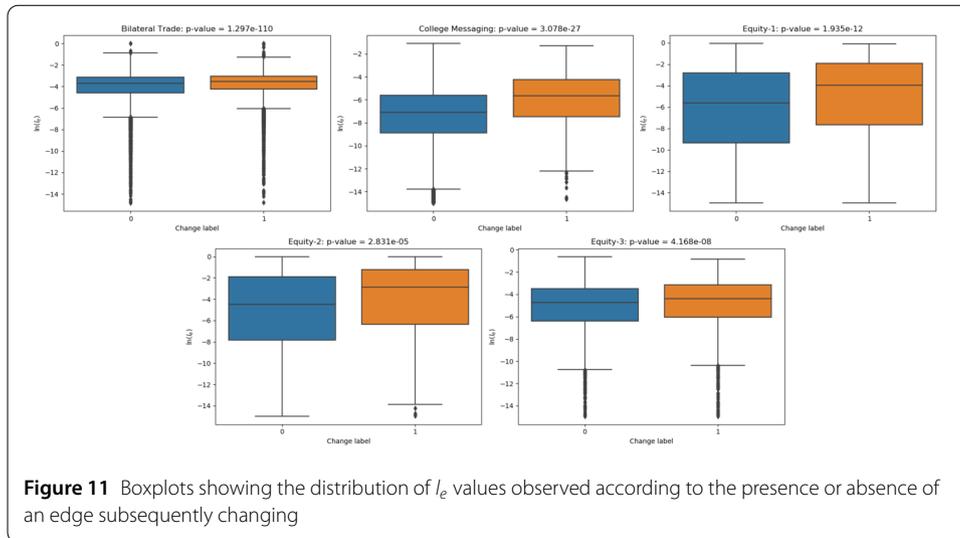
Figure 10 Probability distribution of the values of $\log(l_e)$ for different networks

Instead, we consider the probabilities of observing values of l_e by making use of Kernel Density Estimation to estimate the probability density functions from the data.

Figure 10 shows the estimated Probability Density Functions of the logarithm of the value of l_e . We see from these that for all networks, the values observed for l_e tend to be very small. Omitting the tails of the distributions for diminishingly small values of l_e , we see a similarity in the values of l_e observed across 3 similar equity datasets, and although across all 5 datasets analysed, the distribution is found to be approximately lognormal, the social network shows a much broader distribution of l_e . The peak of the distribution for the college messaging dataset is also much lower, observed at approximately $\ln(l_e) = -8.8$, whereas the bilateral trade dataset shows a peak at -3.3 , and the equity datasets at -3 , -2.5 and -4.2 .

3.2.5 Dynamic observations in real networks

We now address the central concept of the relationship of l_e observed for our real networks and the probability of an edge to change. Figure 11 shows the distributions of the $\ln(l_e)$ values observed for non-changing edges in comparison to changing edges. We see that in all cases, there is a shift in the mean value of $\ln(l_e)$ towards higher values for edges which do change, which would be suggestive of a positive ρ parameter, and potentially the ability to predict the presence of changes given the value of l_e . The smallest shifts are observed for the Bilateral Trade dataset and Equity-3, which show negligible differences in the mean and quartiles of the values of l_e for changes and no changes, suggesting that we might not expect predictability of changes from the values of l_e in these cases. In all



cases, the differences in the mean values of l_e for change vs. no change is significant, with a two-sided t-test showing $p < 0.05$ for all datasets.

To further understand how the value of l_e relates to the probability for edges to change, we look at the distributions of $P(\Delta A_{ij} = 0 | l_e)$ as shown in Fig. 12. Here we see a decreasing probability of $\Delta A_{ij} = 0$ for the bulk of the distribution for increasing l_e for the bilateral trade and Equity-3 datasets, however the rarely observed edges with $l_e > 0.3$ for these datasets show larger probabilities to remain unchanged. We again see a slight initial decrease for Equity-1 and 2 datasets however the relationship is clearly non-linear for large l_e . The college messaging dataset shows a much larger probability in general for edges to remain unchanged, and shows a very slight decrease in probability to remain unchanged for very small l_e values, however is dominated by noise for $l_e > 0.05$.

Referring back to Sect. 3.2, we considered the ideal cases of linear positive, neutral and negative relationships between l_e and the probability of edge changes. In reality, as shown

Table 2 Estimated α and ρ for the 5 real datasets

Dataset	Estimated α	Estimated ρ
Bilateral trade	$0.783 \pm 1.08 \times 10^{-4}$	$0.072 \pm 1.21 \times 10^{-5}$
College messaging	$0.033 \pm 3.03 \times 10^{-6}$	$0.270 \pm 2.76 \times 10^{-6}$
Equity-1	$0.392 \pm 2.74 \times 10^{-4}$	$0.030 \pm 2.82 \times 10^{-5}$
Equity-2	$0.401 \pm 3.13 \times 10^{-5}$	$0.016 \pm 4.93 \times 10^{-5}$
Equity-3	$0.465 \pm 2.57 \times 10^{-4}$	$0.036 \pm 4.21 \times 10^{-5}$

in Fig. 12, we see things are more complex, with different relationships apparent for different l_e ranges. In particular, for edges with lower values of l_e , the negative relationship between the value of l_e and the probability of an edge to remain unchanged suggests that a parameterisation of our model with positive value of ρ would be effective in capturing the behaviour of the bulk of the network. However changes to the small handful of edges with the largest values of l_e are less likely. These observations could suggest that there are a few structurally important edges which act to stabilise a system which would otherwise move towards a regime of instability.

3.2.6 Estimation of α and ρ from data

In Table 2, we present the values of α and ρ estimated for our 5 different datasets. The errors on these estimations are given by the inverse hessian of the Log-Likelihood, which is found by numerical approximation. In comparison with Figs. 5 and 12, we see the ordering of the estimated value of α appears to agree with the positions of the college messaging dataset and the equity datasets. The parameter ρ appears to correspond with the overall gradients observed in Fig. 12 for the bulk of the distributions observed for low values of l_e . These observations suggest that our model is mostly capturing the imbalance of observed changes in the parameter ρ , and the overall average change probability for each dataset in the parameter α .

3.2.7 Edge change predictability

Given the non-zero estimated values of the parameters α and ρ , it is natural to assess the performance of using the value of l_e to predict a subsequent change. Figures 13 and 14 show the Receiver Operating Characteristic and Precision-Recall Curves for the 5 different datasets, and Table 3 shows the corresponding Area Under Curve and balanced accuracy scores.

All datasets are seen to perform slightly better than the dummy model, with better performance seen for the College Messaging dataset and Equity-1 and 2, which also show larger differences in the distribution of l_e across change vs. no change in Fig. 6. Poorer performance is seen for the bilateral trade and Equity-3 datasets, which show similar shaped distributions in Fig. 12 with an initial steep decrease in probability to remain unchanged for increasing l_e , however this trend appears to reverse for $l_e > 0.3$. These datasets also show little difference in the distribution of values observed in Fig. 6 and are found to have low values of ρ . Although the college messaging dataset shows the best performance, particularly in the left hand side of the ROC curve, this is driven by the significant class imbalance with only 5% of the observations showing a non-zero ΔA_{ij} , as opposed to the bilateral trade dataset which shows a 20% proportion of non-zero changes. This is also reflected in the Precision-Recall AUC score for the College Messaging dataset being close to the upper margin of error for the null model.

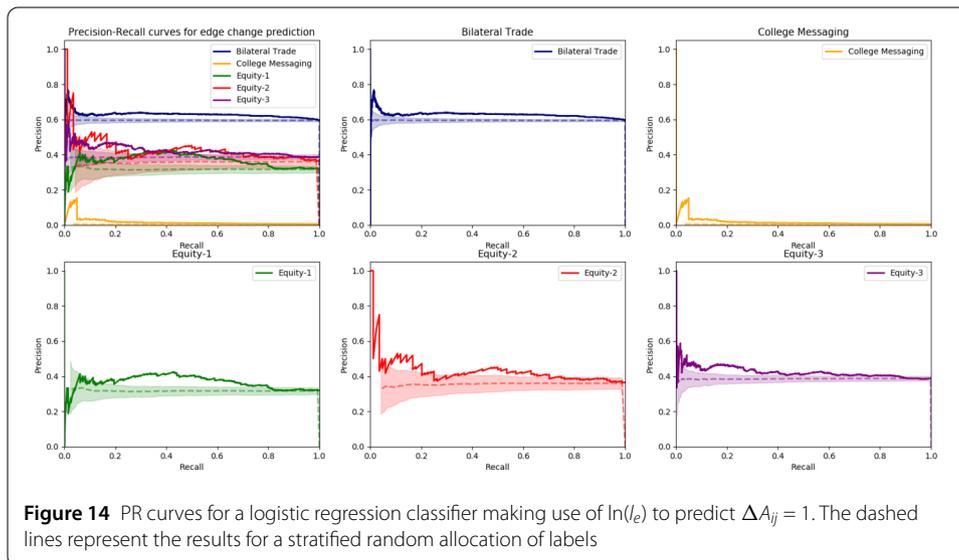
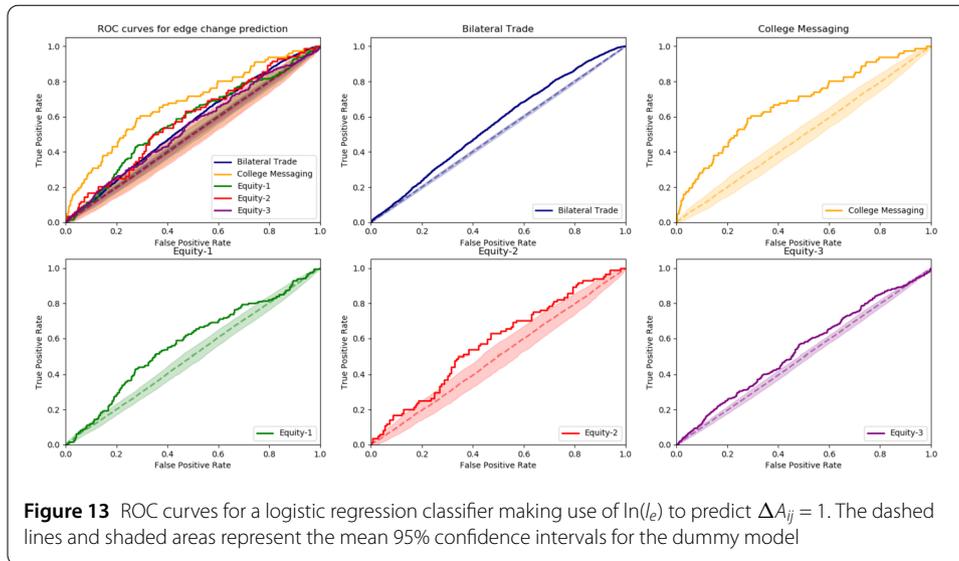
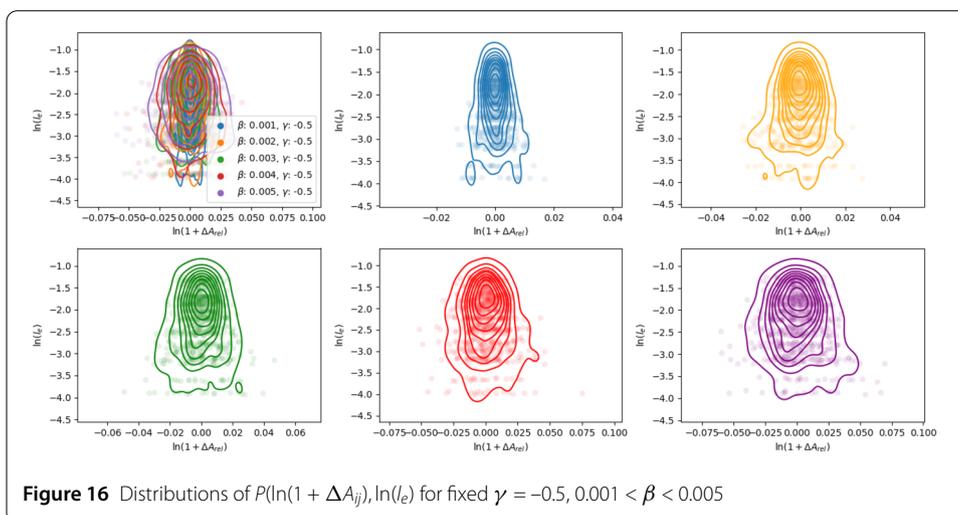
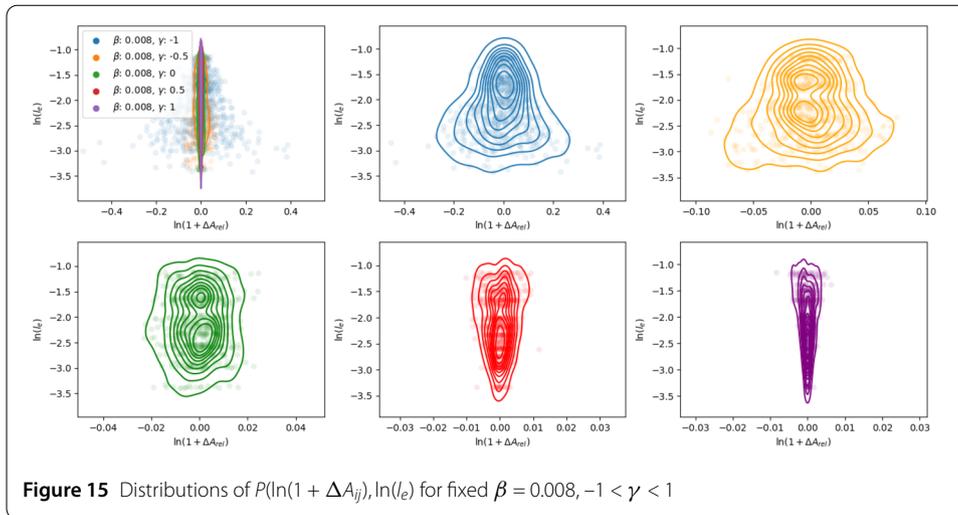


Table 3 Values of Area Under Curve scores for ROC and Precision-Recall curves. Numbers in brackets represent the score achieved by a model which randomly predicts 1 or 0 in proportion to the dataset prior, averaged over 100 trials

Dataset	Balanced accuracy	ROC AUC	Precision-Recall AUC
Bilateral trade	0.542 (0.5±0.0011)	0.554 (0.5±0.0013)	0.628 (0.595±0.0021)
College messaging	0.623 (0.5±0.0027)	0.678 (0.5±0.0029)	0.017 (0.005±0.0045)
Equity-1	0.568 (0.5±0.0039)	0.576 (0.5±0.0046)	0.365 (0.313±0.0062)
Equity-2	0.566 (0.5±0.0061)	0.579 (0.5±0.0073)	0.430 (0.351±0.010)
Equity-3	0.527 (0.5±0.0025)	0.542 (0.5±0.0045)	0.424 (0.381±0.0030)

3.3 Relationship of between l_e and size of weight changes

We now consider if the value of l_e is observed to have an affect on the scale of subsequent edge changes. As in Sect. 3.2, we again consider data generated according to the model in equation (5), and we choose to take $\mathcal{U}_{ij}^t = \mathcal{N}(\mu = 0, \sigma = \beta l_e^{\gamma})$. This introduces two new



parameters, β which controls the width of the distribution of edge changes, and γ which controls the level to which l_e influences the variance of the edge change distribution.

3.3.1 Variation of γ

Figure 15 shows the distributions of $P(\ln(1 + \Delta A_{ij}), l_e)$ for a range of values of γ . We see here that for positive γ , the width of the distribution widens for larger l_e . For negative γ , we see the opposite, that the width of the distribution becomes narrower for larger l_e .

3.3.2 Variation of β

Figure 16 shows the distributions of $P(\ln(1 + \Delta A_{ij}), l_e)$ for a range of values of β . We see here that as β increases, the width of the distributions increase.

3.3.3 Weight distributions for real networks

We now consider the same 5 real datasets considered in Sect. 3.2.5. Figure 17 shows the distributions of $P(\ln(1 + \Delta A_{ij}), \ln(l_e))$ for the case of edges that do change, i.e. $\Delta A_{ij} \neq 0$ for the five real networks. Note that ΔA_{ij} refers to the relative change in the value of the edge

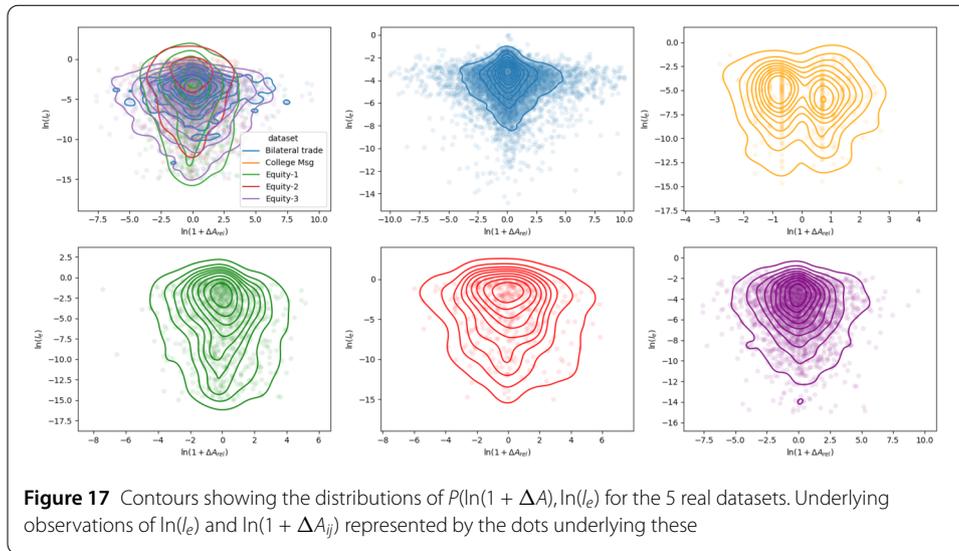


Table 4 Estimated β and γ for the 5 real different datasets

Dataset	Estimated β	Estimated γ
Bilateral trade	$1.11 \times 10^{-4} \pm 8.93 \times 10^{-18}$	$1.93 \pm 1.10 \times 10^{-27}$
College messaging	$7.77 \times 10^{-5} \pm 1.10 \times 10^{-16}$	$1.15 \pm 3.07 \times 10^{-27}$
Equity-1	$1.51 \times 10^{-5} \pm 4.55 \times 10^{-24}$	$1.35 \pm 2.55 \times 10^{-36}$
Equity-2	$1.42 \times 10^{-4} \pm 1.36 \times 10^{-26}$	$1.32 \pm 5.81 \times 10^{-37}$
Equity-3	$1.63 \times 10^{-4} \pm 2.20 \times 10^{-26}$	$1.42 \pm 1.43 \times 10^{-36}$

weight from t_0 to t_1 , which takes values in the interval $[-1, \infty]$, and l_e is measured at time t_0 . Infinite values for ΔA_{ij} , corresponding to the case of a new edge appearing, were observed but are not captured in the plots. The prominence of these across the different datasets are 4.7% of the bilateral trade dataset, 0.086% of the college messaging dataset, 0.012%, 0% and 0.0028% of the equity datasets.⁴ We see a slight widening of the distributions for larger values of l_e for Equity-1 and 2 datasets, and to a larger extent for the third equity dataset. The bilateral trade dataset shows initial widening as l_e increases, however narrows again for the largest l_e edges. The college messaging dataset shows two distinct peaks, corresponding to changes in edge weight of ± 1 , which are over-represented in this dataset as it is unweighted, and the edge weight solely represents the count of interactions in the time window of consideration. The slight widening for larger l_e for all datasets is suggestive of a positive relationship between the value of l_e and the variance of the distribution of subsequent edge changes.

3.3.4 Parameter estimation for β and γ

The estimated values of the parameters β and γ for the different datasets are shown in Table 4. All 5 datasets show positive values of γ , suggestive of a relationship between the width of the distribution of edge changes and the value of l_e . The dataset with the highest values for γ , the bilateral trade dataset, also shows the largest level of bias towards larger change distribution width for higher l_e in Fig. 17. Correspondingly, the lowest γ value is seen for the college messaging dataset, which shows the least bias towards larger changes

⁴The prominence of new edges has been significantly reduced by focusing on the giant component.

occurring for larger values of l_e . The values for β are similar across the 5 datasets, and all relatively low. It is difficult to draw conclusions from these, as the behaviours controlled by the two parameters cannot be separated and observed alone in the distributions in Fig. 17.

4 Discussion & conclusion

The ability to understand how microscopic changes in networks affect the macroscopic evolution across time is one of the key challenges in dynamic network analysis. In this study we have begun to explore the use of derivatives of network spectra to capture this. We derive a measure of edge based structural influence, l_e , and explore the extent to which the value is indicative of future changes. We first of all demonstrated that for small and isolated perturbations applied to the network, the eigenvalue derivative is approximated well by equation (2). However, we observe the approximation breaks down for multiple changes happening during the same time snapshot, suggesting that the measure may be more suited to a continuous or pseudo-continuous representation of network evolution, in which each time snapshot contains a single edge change.

Considering the 5 real datasets, we observe lognormal distributions of the values of l_e , indicating structural influence dominated by a small handful of edges. We propose a model in which the probability for an edge to change is given by αl_e^ρ . This model allows us to control the extent to which l_e dictates the propensity for an edge to change, and also controls the scale of a subsequent change. Focusing on the former, we observe similarities in the shapes of the distributions of $P(\Delta A_{ij} = 0, \ln(l_e))$ when generating synthetic networks according to this model and those observed in the data, and the values observed for α and ρ are suggestive of a relationship between the value of l_e and the subsequent presence of change. In using l_e in a logistic regression classifier to predict change, we see that l_e is slightly predictive of change in all cases, but only marginally so for the case of the bilateral trade and Equity-3 datasets. This corresponds with our observations of small values of ρ for these datasets, along with similar, non-linear distributions shapes for the probability of no change for increasing l_e . These observations indicate that the static structural importance can be indicative of the presence of a subsequent change, however more work is needed to understand the shape of the distribution and the identification of different l_e regimes. We will also consider taking a similar approach with other measures of edge importance, for example edge gravity [18]. More work is also needed to understand the subsequent impact on the global network structure of an edge changing. It may be that a change to an influential edge could act to destabilise a system; conversely, the change could move the system towards a state of stability. We will look to investigate this in future analyses.

We note here that α and ρ themselves are useful parameters that could be used to classify networks according to their growth stability. A large value of α would be an indicator for larger levels of overall network activity. A network with very large ρ would be characterised by changes occurring to the edges with the largest l_e , conversely, a network with very small ρ would see changes distributed across all edges, regardless of the value of l_e . In the context of financial markets, these contrasting situations would require different approaches, and ρ could be used by policy makers to inform which asset classes should be monitored as a whole (for the case of small ρ) or following an approach targeting those edges with the highest l_e .

Our model doesn't account for edges appearing and disappearing in the network, and assumes that edge changes are independent of each other. For the first limitation, we note

that edge appearance and disappearance would be unlikely to heavily influence the behaviour of the Equity networks, as we observed very low percentages (0.012%, 0% and 0.0028%) of new edges appearing,⁵ but for the other two networks this behaviour is much more prominent at 4.7% for the bilateral trade network and 0.086% for the college messaging network. The measure l_e itself is able to assess the importance of an edge that subsequently disappears, and also those that appear between two existing nodes, so understanding how these appearances and disappearances can be captured in a model for network growth would be highly beneficial for future work. On the second point, we noted in our exploration of toy networks that the ability of our approximation of the eigenvalue derivative breaks down for multiple edge changes present. Conversely, many works such as Bandi et al. have noted that predictability is aggregation scale specific. In future work we will thus investigate the trade-off between improved approximation of l_e for the quasi-continuous limit in which each time snapshot contains a single edge change, and improved predictability for larger aggregation scales. In addition to this, further analysis is needed to assess the effectiveness of l_e as an indicator for risk, as so far we understand that the value of l_e bears some relationship to how the network subsequently changes, but we have not yet considered the resultant changes of edges with high values of l_e , and how these have an effect on the rest of the network in terms of risk and stability. This is another area we will pursue in future work. We will also consider extending our methods to consider structural node importance, which is of use to policy makers who may wish to monitor which players could have an adverse impact on markets. It is also worth noting that although using raw transaction data gives us the lowest granularity view of the data, our work has so far not considered the higher order effects of trading behaviour on price. Such an effect results in the influence of edges reaching disconnected components, which cannot be captured by our methods, so we will consider generalising our methods to allow for networks with disconnected components. Finally, we will consider using our methods for classification of a large number of networks, and also extend our methods to understand the parameters which control the resultant weight changes.

Supplementary information

Supplementary information accompanies this paper at <https://doi.org/10.1140/epjds/s13688-021-00279-6>.

Additional file 1. This document contains Supplementary Information, including addition results, derivations and analysis of non-public datasets (PDF 2.0 MB)

Acknowledgements

The authors do not wish to acknowledge any additional support in this research.

Funding

This research received no external funding.

Availability of data and materials

The datasets referred to as Equity-1, Equity-2 and Equity-3 in this paper, were extracted from a dataset of transaction reports collected by the FCA under MIFID II regulations. The datasets were used under agreement from the data owners at the Financial Conduct Authority for the current study, and are not publicly available. A data note describing the data is shown in the Additional file 1, and comparable analysis is conducted for all investigations for two open source datasets (see below). The dataset referred to as Bilateral Trade are hosted by Katherine Barbieri, University of South Carolina, and Omar Keshk, Ohio State University, available at <http://correlatesofwar.org>. The dataset referred to as College Messaging

⁵This is largely due to the preprocessing steps applied to the network, since we are only considering the giant component.

are available in the Stanford Large Network Dataset Collection repository, <https://snap.stanford.edu/data/CollegeMsg.html>. An implementation of the methods referenced in this paper can be found at [50].

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

Author IS conducted the research and authored the final manuscript with support and advice given by Authors FC and PB. All authors approved the final manuscript.

Author details

¹Financial Conduct Authority, Endeavour Square, London, United Kingdom. ²Department of Computer Science, University College London, London, United Kingdom. ³Systemic Risk Centre, London School of Economics, London, United Kingdom. ⁴London Mathematical Laboratory, London, United Kingdom.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 7 January 2021 Accepted: 3 May 2021 Published online: 08 May 2021

References

1. Nier E, Yang J, Yorulmazer T, Alentorn A (2007) Network models and financial stability. *J Econ Dyn Control* 31:2033–2060
2. Silva TC, Guerra SM, Tabak BM, de Castro Miranda RC (2016) Financial networks, bank efficiency and risk-taking. *J Financ Stab* 25:247–257. <https://doi.org/10.1016/j.jfs.2016.04.004>
3. Allen HJ (2017) Financial stability regulation as indirect investor/consumer protection regulation: implications for regulatory mandates and structure. *Legal Studies Research Paper Series Research Paper* 17-5
4. Haldane A, May R (2011) Systemic risk in banking ecosystems. *Nature* 469:351–355. <https://doi.org/10.1038/nature09659>
5. Ren X, Lu L (2014) Review of ranking nodes in complex networks. *Chin Sci Bull* 59:1175. <https://doi.org/10.1360/972013-1280>
6. Stegehuis C, van der Hofstad R, van Leeuwen SH (2016) Epidemic spreading on complex networks with community structures. *Sci Rep* 6:29748. <https://doi.org/10.1038/srep29748>
7. Ren G, Wang X (2014) Epidemic spreading in time-varying community networks. *Chaos, Interdiscip J Nonlinear Sci* 24(2):023116
8. Lai Y, Motter A, Nishikawa T (2004) Attacks and cascades in complex networks. *Lect Notes Phys* 650:299–310. <https://doi.org/10.1007/978-3-540-44485-514>
9. Xu H, Zhang J, Yang J, Lun L (2018) Identifying important nodes in complex networks based on multiattribute evaluation. *Math Probl Eng*. <https://doi.org/10.1155/2018/8268436>
10. Moreno Y, Nekovee M, Pacheco AF (2004) Dynamics of rumor spreading in complex networks. *Phys Rev E* 69:066130. <https://doi.org/10.1103/PhysRevE.69.066130>
11. Restrepo JG, Ott E, Hunt BR (2005) Weighted percolation on directed networks. *Phys Rev Lett* E 71:036151
12. Bollobás B, Borgs C, Chayes J, Riordan O (2010) Percolation on dense graph sequences. *Ann Probab* 38(1):150–183
13. Wang Y, Di Z, Fan Y (2011) Identifying and characterizing nodes important to community structure using the spectrum of the graph. *PLoS ONE* 6. <https://doi.org/10.1371/journal.pone.0027418>
14. Lü L, Pan L, Zhou T, Zhang Y, Stanley HE (2015) Toward link predictability of complex networks. *Proc Natl Acad Sci* 112(8):2325–2330. <https://doi.org/10.1073/pnas.1424644112>
15. Restrepo J, Ott E, Hunt B (2006) Characterizing the dynamical importance of network nodes and links. *Phys Rev Lett* 97:094102. <https://doi.org/10.1103/PhysRevLett.97.094102>
16. Battiston S, Puliga M, Kaushik R, Tasca P, Caldarelli G (2012) Debrank: too central to fail? financial networks, the fed and systemic risk. *Sci Rep* 2:541. <https://doi.org/10.1038/srep00541>
17. Barucca P, Lillo F (2018) The organization of the interbank network and how ecb unconventional measures affected the e-mid overnight market. *Comput Manag Sci* 15. <https://doi.org/10.1007/s10287-017-0293-6>
18. Helander ME, McAllister S (2018) The gravity of an edge. *Appl Netw Sci* 3. <https://doi.org/10.1007/s41109-018-0063-6>
19. Yu E, Chen D, Zhao J (2018) Identifying critical edges in complex networks. *Sci Rep* 8:14469. <https://doi.org/10.1038/s41598-018-32631-8>
20. Wang Y, Chakrabarti D, Wang C, Faloutsos C (2003) Epidemic spreading in real networks: an eigenvalue viewpoint. In: 22nd international symposium on reliable distributed systems. Proceedings
21. Pei S, Makse HA (2013) Spreading dynamics in complex networks. *J Stat Mech Theory Exp*
22. Bardoscia M, Bianconi G, Ferrara G (2019) Multiplex network analysis of the uk over-the-counter derivatives market. *Int J Econ Finance* 24(4):1520–1544. <https://doi.org/10.1002/ijfe.1745>
23. May RM (1972) Will a large complex system be stable? *Nature* 238
24. Restoy F (2019) Market integration: the role of regulation. Speech by Chairman of the Financial Stability Institute, Bank for International Settlements, at the IIF Market fragmentation roundtable, Washington DC, United States
25. Samuelson PA (1967) General proof that diversification pays. *J Financ Quant Anal* 2:1–13
26. Bardoscia M, Battiston S, Caccioli F, Caldarelli G (2016) Pathways towards instability in financial networks. *Nat Commun* 8:14416. <https://doi.org/10.1038/ncomms14416>
27. Markose S, Giansante S, Gatkowski M, Shaghghi AR (2010) Too interconnected to fail: financial contagion and systemic risk in network model of cds and other credit enhancement obligations of us banks. *Comisef Working Papers Series WPS-033* (2010)

28. Caccioli F, Marsili M, Vivo P (2009) Eroding market stability by proliferation of financial instruments. *Eur Phys J B, Condens Matter Complex Syst* 71(4):467–479
29. Tobin J (1978) The new economics one decade older. The Eliot Janeway lectures on historical economics in honour of Joseph Schumpeter 1972 (Princeton University Press, Princeton, US). *East Econ J* IV:153–159
30. Bianconi G, Galla T, Marsili M, Pin P (2009) Effects of tobin taxes in minority game markets. *J Econ Behav Organ* 70
31. Brock WA, Hommes CH, Wagener FOO (2009) More hedging instruments may destabilize markets. *J Econ Dyn Control* 33:1912–1928
32. Aste T, Shaw W, Di Matteo T (2010) Correlation structure and dynamics in volatile markets. *New J Phys* 12:085009. <https://doi.org/10.1088/1367-2630/12/8/085009>
33. Mazzarisi P, Barucca P, Lillo F, Tantari D (2020) A dynamic network model with persistent links and node-specific latent variables, with an application to the interbank market. *Eur J Oper Res* 281:50–65
34. Barucca P, Lillo F (2015) Disentangling bipartite and core-periphery structure in financial networks. *Chaos Solitons Fractals* 88. <https://doi.org/10.1016/j.chaos.2016.02.004>
35. Barucca B, Bardoscia M, Caccioli F, D'Errico M, Visentin G, Battiston S, Caldarelli G (2020) Network valuation in financial systems. *Math Finance* 30:1181–1204
36. Soramäki K, Bech M, Arnold J, Glass R, Beyeler W (2007) The topology of interbank payment flows. *Physica A* 379:317–333
37. Barabási A-L, Albert R (1999) Emergence of scaling in random networks. *Science* 286(5439):509–512. <https://doi.org/10.1126/science.286.5439.509>
38. Falkenberg M, Lee J, Amano S, Ogawa K, Yano K, Miyake Y, Evans T, Christensen K (2020) Identifying time dependence in network growth. *Phys Rev Res* 2:023352
39. Bazzi M, Jeub LGS, Arenas A, Howison SD, Porter MA (2020) A framework for the construction of generative models for mesoscale structure in multilayer networks. *Phys Rev Res* 2:023100. <https://doi.org/10.1103/PhysRevResearch.2.023100>
40. Peixoto TP, Rosvall M (2017) Modelling sequences and temporal networks with dynamic community structures. *Nat Commun* 8. <https://doi.org/10.1038/s41467-017-00148-9>
41. Watts D, Strogatz S (1998) Collective dynamics of 'small-world' networks. *Nature* 393:440–442. <https://doi.org/10.1038/30918>
42. Bianconi G, Barabási AL (2001) Competition and multiscaling in evolving networks. *Europhys Lett* 54:436–442
43. Caldarelli G, Capocci A, De Los Rios P, Muñoz MA (2002) Scale-free networks from varying vertex intrinsic fitness. *Phys Rev Lett* 89:258702. <https://doi.org/10.1103/PhysRevLett.89.258702>
44. Kobayashi T, Géniois M (2020) Two types of densification scaling in the evolution of temporal networks. *Phys Rev E* 102
45. Holme P, Saramaki J (2012) Temporal networks. *Phys Rep* 519:97–125
46. Tang J, Musolesi M, Mascolo C, Latora V, Nicosia V (2010) Analysing information flows and key mediators through temporal centrality metrics. In: SNS'10: proceedings of the 3rd workshop on social network systems, pp 1–6. <https://doi.org/10.1145/1852658.1852661>
47. Kuhn HW, Tucker AW (1950) Nonlinear programming. In: Proceedings of the second Berkeley symposium on mathematical statistics and probability, pp 481–492
48. Barbieri K, Keshk O (2016) Correlates of War Project Trade Data Set Codebook, Version 4.0. Online: <http://correlatesofwar.org>
49. Panzarasa P, Opsahl T, Carley KM (2009) Patterns and dynamics of users' behavior and interaction: network analysis of an online community. *J Am Soc Inf Sci Technol* 60(5):911–932
50. Seabrook I (2020) Structural edge importance—python implementation. Github repository

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)
