

TIME SERIES APPROACH TO THE EVOLUTION OF NETWORKS: PREDICTION AND ESTIMATION

ANNA BYKHOVSKAYA

ABSTRACT. The paper analyzes non-negative multivariate time series which we interpret as weighted networks. We introduce a model where each coordinate of the time series represents a given edge across time. The number of time periods is treated as large compared to the size of the network. The model specifies the temporal evolution of a weighted network that combines classical autoregression with non-negativity, a positive probability of vanishing, and peer effect interactions between weights assigned to edges in the process. The main results provide criteria for stationarity vs. explosiveness of the network evolution process and techniques for estimation of the parameters of the model and for prediction of its future values.

Natural applications arise in networks of fixed number of agents, such as countries, large corporations, or small social communities. The paper provides an empirical implementation of the approach to monthly trade data in European Union. Overall, the results confirm that incorporating non-negativity of dependent variables into the model matters and incorporating peer effects leads to the improved prediction power.

KEYWORDS: multivariate time series, non-negative time series, networks, dynamic tobit, LAD.

Date: November 20, 2020.

The author would like to thank Donald Andrews, Vadim Gorin, Bruce Hansen, Peter Phillips, Jack Porter, and Larry Samuelson for valuable comments, suggestions, and encouragement throughout the duration of the project.

1. Introduction

1.1. Motivation. This paper studies non-negative time series. We are especially interested in multivariate time series, whose components can be interpreted as edges of some network (if a network is composed of only two vertices, then we are dealing with a basic one-dimensional time series). The novelty of our view on networks lies in the time series approach. Citing Diebold and Yilmaz (2014), “the network and multivariate time series literatures have much to learn from each other”, and our paper makes a step in this direction. We develop a model and a consistent estimator of its parameters. The three main features of our model are non-negativity of weights associated with edges, positive probability of vanishing of each edge, and possibility for edges to affect each other in the multivariate setting.

Non-negative time series often arise in different socio-economic settings. For example, production decisions by firms cannot be negative. Similarly, technology adoption by firms or countries is a non-negative variable, which evolves over time. Macroeconomic variables such as exchange rates, T-bill rates and so on are non-negative by definition. Finally, if we look at the number of passengers travelling from country A to country B at a given day, it is again either zero or positive.

One can note that in almost all of the above examples zero occurs with positive probability: there are times when firms decide not to produce at all, or two countries do not sell some good to each other, or a group of people do not call each other and so on. Thus, if one wants to find a suitable model for such non-negative processes, allowing it to take zero values with positive probability is an important condition. However, this feature cannot be captured by linear models or by their simplest modifications such as GARCH.

Moreover, we look not only at a one-dimensional time series, but also at multivariate ones. For instance, instead of considering the traffic flow from one given country to another, we can track the whole set of countries, and get a non-negative weighted network, which governs the amount of travel. In the same spirit, we can construct a time-varying network of trade between different countries or firms or a social network, which measures the amount of communication between people. In this paper we focus on networks (or weighted graphs) as a main example of a non-negative structure. In contrast to other cases, networks may

potentially involve interactions between the coordinates corresponding to their edges. For example, if a firm A is in search of a provider of some intermediate good, it may contact firm B, with which it has well-established trading relationship. Even if B cannot provide such intermediate good, it may recommend firm C, with which B in turn has a well-established relationship. Robinson and Stuart (2006) illustrate that the whole network of past alliances in the biotechnology industry affects the structure and size of alliance agreements between any two given firms in the industry in the future. (For more examples on the role of social and economic networks in real life see, for instance, Jackson (2010).)

Often one may need to be able to predict a structure of a network as it evolves. It may be important to know whom is better to target with ads, election campaigns, or other type of information. Countries may need to know the future trading patterns to decide production of which types of goods to support. We also refer to Holme and Saramäki (2012) for an interdisciplinary review highlighting the importance of the temporal structure of networks. The availability of time series data for social and economic networks, such as call detail records and international and financial trade data has grown tremendously. In this paper we use such time-series data to estimate parameters of network evolution and to make predictions. When one goes to the data, working with networks as a multivariate time series (when the number of time periods is large compared to the size of the network) and not as a cross-section or panel makes a difference between our paper and what is currently done in networks literature (see e.g., Bramoullé et al., eds (2016)). One natural example of applicability of our setting is when agents are countries or large companies (as there are not so many of them). For social communities this is also of relevance, as the dynamics of small groups of people might be very different from the large ones (see e.g., Palla et al. (2007)).

1.2. Results. We deal with multivariate non-negative time series, which we interpret as weighted networks (although alternative interpretations and applications are also possible and some of them are presented in Section 7). Each edge is associated with a number, which evolves with time. We want to capture three essential aspects of networks. Those properties are non-negativity of weights, possibility not to have an edge between any two nodes with positive probability, and possibility for the past of the network to affect all edges today.

These three features lead us to modelling the evolution of a network as a non-linear process

$$(1) \quad y_{ijt} = [\alpha_{ij} + \beta_{ij}y_{ij,t-1} + \gamma_{ij}z_{ij,t-1} + u_{ijt}]_+,$$

where y_{ijt} is the weight of the edge ($i \rightarrow j$), $[\cdot]_+$ stands for positive part, α_{ij} , β_{ij} , γ_{ij} are unknown coefficients, which need to be estimated, u_{ijt} is a random error, which is independent across time t , but may be correlated along edges (i, j) , and $z_{ij,t-1}$ is a peer effect or an interaction term. That is, $z_{ij,t-1}$ is a function (assumed to be explicitly known) of past periods of the network. It allows edges to affect each other in the future. Note that we do not impose linearity on z . It may be a non-linear function giving us a lot of flexibility in how we model peer-effects. Moreover, the positive part per se creates nonlinearity, which leads to technical difficulties. The equation (1) can be obtained as a solution to a certain utility maximization problem, as we outline in Section 2. One can add more regressors to Eq. (1) and most of our results continue to hold in the extended setting (see Remarks 4 and 5). We note that there is a large number of much more sophisticated models for the network evolution in the literature (e.g., Pin and Rogers (2015) where agents play a prisoner's dilemma every period and choose connections based on the observed outcomes), yet the estimation procedures become more advanced and model-specific with each layer of complexity. It becomes preferable to address the basic case, so that the general principles and challenges can be identified.

Our setting allows for arbitrary heteroskedasticity in the errors u_{ijt} , that is, arbitrary correlation across individuals. This differs a lot from what is feasible when one relies on cross-sectional variation instead of the variation across time. The classical approach to networks (n is large and $T = 1$) assumes either dyadic error structure where $Cov(u_{ijt}, u_{klt}) = 0$ if $i \neq k, \ell$ and $j \neq k, \ell$ (e.g., Graham (2017), Graham (2019)) or other type of decreasing correlation, so that one can average over individuals and use law of large numbers and central limit theorem.

In principle, an alternative basic approach could be to consider a latent censored model, see e.g., Wei (1999) and references therein. That is to assume that there is some unobserved underlying process y_t , which is allowed to take negative values. The researcher instead

observes the truncated process of the form $y_t^* = y_t \mathbf{1}(y_t > 0)$. In our case there need not exist such an underlying process. Moreover, the censored structure makes the observed process non-Markovian. Thus, it is hard to make predictions. Our approach, in contrast, guarantees, a Markovian structure (in a state space taking into account several time lags). So it is easy to make predictions based on our model.

Let us describe our main findings. First, we show a sufficient condition for the model to be stationary. The model does not belong to any known class and requires a different treatment. We normalize z_{ijt} in a way which insures that peer effects do not grow faster than their maximal argument (see Section 2.3 for more details). We prove that if $\max(0, \beta_{ij}) + |\gamma_{ij}|$ is uniformly over i, j bounded by some constant smaller than one, then the process is strongly mixing and converges (when started from any initial conditions) to a stationary distribution as $t \rightarrow \infty$. For the special case of the model where $\gamma_{ij} \equiv 0$, we provide the full classification of asymptotic behavior of the process: the change from stationary to explosive behavior is at $\alpha_{ij} = 0, \beta_{ij} = 1$. At that boundary case, the process rescaled by \sqrt{T} converges to the absolute value of a Brownian motion as $T \rightarrow \infty$.¹ One distinction with the usual linear autoregression model is that the $\alpha_{ij} < 0, \beta_{ij} = 1$ case is stationary due to the presence of positive part.

We remark that the non-linearities created by $[\cdot]_+$ and by z_{ijt} make asymptotic analysis challenging. The arbitrary peer-effect function $z_{ij}(\cdot)$ significantly complicates the setting, leading to a potentially non-Lipschitz dependence on the past. For example, the triangular peer-effect function $z_{ijt} = \sum_k \frac{\sqrt{y_{ikt}y_{kjt}}}{n-2}$, which we use in the empirical example, is not Lipschitz. Thus, the classical arguments for stationarity, which are based on the contraction mapping, are not applicable and one has to come up with a different approach. To overcome those problems, we develop a proof based on the large deviations principle for $\{y_{ijt}\}$ and analyze the expected time it takes for the process to jump to zero. Then we apply renewal theorem to obtain mixing and stationarity independent of initial conditions. As far as we know, our stationarity results are new; the only relevant papers seems to be de Jong and Herrera (2011), Hahn and Kuersteiner (2010), and Michel and de Jong (2018). Among other results,

¹For an earlier proof of an equivalent result in the boundary case see Cavaliere (2004).

in these papers sufficient conditions for the existence of a stationary solution are found for the special case of Eq. (1) dealing with a one-dimensional process with z linearly depending on the past periods. This particular stationary solution (rather than a general one) is further shown to be strongly mixing.

Let us stress that in contrast to classical autoregression, where a continuous distribution of the errors is important to establish strong mixing (see Withers (1981) and Andrews (1984) for examples of non-strongly mixing AR processes with discrete errors), in our setting we only need the errors to have support that is unbounded from below. Thus, the distribution is not required to be continuous. This also differs from Hahn and Kuersteiner (2010), who require a continuous distribution of the errors to obtain strong mixing.

Second, we discuss how to estimate the parameters of the model. The nonlinearities due to censoring and peer effects significantly complicate estimation. The ordinary least squares (OLS) estimator turns out to be inconsistent in our setting. This matches a similar inconsistency for censored regression models, c.f. discussion at the end of Section 4.2 in Amemiya (1984). Our approach builds on and extends the least absolute deviations (LAD) method, which was used in the context of censored regression in Powell (1984). We obtain an estimator by minimizing with respect to α_{ij} , β_{ij} , γ_{ij} the sum of the absolute differences between y_{ijt} and $[\alpha_{ij} + \beta_{ij}y_{ijt-1} + \gamma_{ij}z_{ijt-1}]_+$. Theorems from Powell (1984) are not applicable in our setting because the independence assumptions from that paper do not hold. Thus, novel ideas are required.

Some of our results cover explosive cases and we show that while the large T asymptotics of the system changes drastically, the LAD estimator is still consistent. To our knowledge, we are the first to analyze LAD in the explosive setting. Our findings are in line with results on the consistency of OLS in the explosive autoregressive model (see White (1958) and Anderson (1959) for the model without a constant and Wang and Yu (2015) for the model with an intercept.)

We remark that, in general, minimization of absolute deviations in models with a positive part is a non-convex problem and designing numerical algorithms requires special care (see e.g., Khan and Powell (2001)). Similarly, asymptotic normality in the censored cross-section

model in Powell (1984) relies on certain continuity properties of a function of the true parameter value and the error distribution, which are hard to check. In contrast, we find that whenever z has non-negative support and the true α_{ij} , β_{ij} , γ_{ij} are all positive, neither of these problems exist in our setting: the optimization problem is convex, asymptotic normality does not rely on any additional conditions, and the asymptotic variance of the LAD estimator is given by a simple formula.

We also show in the Supplementary Material how to correct the OLS procedure to restore consistency, yet in doing so one needs to ignore a lot of observations, and, thus, the accuracy of the estimation decreases significantly. Further, if we assume the errors u_{ijt} are Gaussian, then we can explicitly write down the likelihood function. We prove that in this case the maximum likelihood estimator (MLE) is consistent. However, for the semiparametric case when the distribution of u_{ijt} is not specified, there is no guarantee that the Gaussian MLE is consistent.

Being able to consistently estimate γ_{ij} allows us to analyze the importance of peer effects in the network formation process. This relates our paper to the vast literature on estimating peer effects in various network models (e.g., see Blume et al. (2011) for review). Several approaches rely on specific functional forms to ensure a consistent estimate (e.g., triangular peer effects in Graham (2016), which resemble our triangles example after Assumption 3 in Section 2.3), while others introduce model free notions for the strength of peer effects (such as connectedness in Diebold and Yilmaz (2015)). However, we have not seen in the literature estimation techniques based on variation across time in non-linear Markov evolution model, as in our work.

Finally, we propose an objective quality measurement of the model by looking at the absolute value of the prediction error at time $t + T'$ for the model estimated from the data from time t and up to time $t + T' - 1$ for all $t = 0, \dots, T - T'$, where T' is the size of the window used for estimation. The choice of absolute values instead of the more standard squared differences is motivated by the fact that when α_{ij} , β_{ij} , γ_{ij} are exactly known, the optimal 1-step-ahead forecast in L_1 does not depend on the density of u , while this observation does not hold in L_2 norm (see Section 5 for more details).

To illustrate our methodology, we apply the framework to monthly trade data between European Union countries for pharmaceutical products. We use the basic prediction “tomorrow=today” as a benchmark for comparing the prediction power, and the interactions (peer effects) are modelled by a function based on the “friend of my friend is my friend” principle. In our experiments any model-based estimation techniques lead to an improvement of the power over benchmark case. The basic linear OLS procedure, which ignores the positivity of weights, leads to the worst results among model-based estimations. The MLE estimator performs better, and the LAD estimator leads to the best results. The addition of peer effects z_{ijt} (vs. setting $\gamma_{ij} = 0$ in the model specification) also leads to improved prediction power. Results of the Diebold-Mariano test (Diebold and Mariano (1995)) and the tests for the significance of peer effects, which are reported in Section 6, support the above conclusions.

1.3. Outline of the paper. Section 2 presents the model and the main equation of interest. All assumptions are stated in that section, as well as sufficient conditions for stationarity. Section 3 discusses the special case of the model when there are no peer effects, so that evolution of each edge is a separate process. The full classification in terms of stationary/explosive behavior is established in this case. Section 4 discusses estimation of the model, while Section 5 proposes the method to measure predictive power. Section 6 applies the model to the trade of pharmaceutical products in European Union. Section 7 discusses extensions of our setting. Finally, Section 8 concludes. All proofs, unless otherwise noted, are in the Supplementary Material, which as well contains results on OLS and MLE estimation and additional lemmas used to prove the main theorems.

2. Model

2.1. Set up. We analyze a multivariate time series, which we interpret as a network. The network consists of n vertices and evolves across time. The network is observed over T periods. In our asymptotic results we assume that n is fixed and T goes to infinity. Thus, we are dealing with many observations of an evolving small network.

Our model allows both for undirected and directed networks. The main example of the former is a social weighted network, where nodes represent people and edges represent how

much time they spend together. E.g., one can use phone call data as a measure of friendship (the more two people text or talk to each other, the closer their relationship is). For the case of a directed network, the applications are mostly for firms or countries and trade between them. Looking separately at exports and imports, we get a directed network.

The equation of interest is

$$(2) \quad y_{ijt} = \left[\alpha_{ij} + \beta_{ij}y_{ijt-1} + \sum_{\kappa=1}^K \gamma_{ij}^{\kappa} p_{ij}^{\kappa} \left(\{y_{kls}\}_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} \right) + u_{ijt} \right]_+,$$

where u_{ijt} is a random error, $t = 1, \dots, T$ stands for time, and $i, j = 1, \dots, n$ stand for agents (in the undirected case $i < j$, in the directed case $i \neq j$). So that y_{ijt} can be interpreted as either how much i and j talk at time t or the amount of trade from i to j .

In Eq. (2) $\alpha_{ij} \in \mathbb{R}$ is a proxy for homophily, i.e. how similar i and j are. The larger α_{ij} is, the stronger is the link connecting i and j . We allow β_{ij} and γ_{ij}^{κ} , $\kappa = 1, \dots, K$ to be of any sign. The coefficient β_{ij} measures the dependence on the own past. The larger β_{ij} is, the more the link between i and j yesterday affects its weight today. Coefficients γ_{ij}^{κ} capture the dependence on the peer effects/interactions p_{ij}^{κ} . Functions p_{ij}^{κ} serve as various aggregators of the past structure of the network in a way that affects the current state. Dependence of the peer-effect functions on more than one past period captures the fact that it may take multiple periods for the information about others to be transmitted. The multitude of peer-effect terms gives a lot of flexibility and allows one to capture very general, diverse forms of interactions. A lot of features of networks can be captured through p_{ij}^{κ} . Finally, the positive part in Eq. (2) creates nonlinearity and leads to a positive mass at zero.

The model is initialized at $t = 1 - H, \dots, 0$ by arbitrary values (possibly random). To be more precise, we assume that as T goes to infinity, H does not grow and $y_{ijt} = O(1)$ for $t = 1 - H, \dots, 0$.

Remark 1. We can incorporate more lags of the autoregressive component. E.g., by adding extra peer-effects $\sum_{\tau=2}^P \beta_{ij}^{\tau} y_{ijt-\tau}$, we get $\sum_{\tau=1}^P \beta_{ij}^{\tau} y_{ijt-\tau}$ instead of $\beta_{ij} y_{ijt-1}$.

For simplicity, in the rest of the paper we focus on the special case of Eq. (2) where $K = 1$,

$$(3) \quad y_{ijt} = \left[\alpha_{ij} + \beta_{ij}y_{ijt-1} + \gamma_{ij}p_{ij} \left(\{y_{kls}\}_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} \right) + u_{ijt} \right]_+.$$

Focusing on $K = 1$ allows us to shorten the notation. As noted in the Remarks 4 and 5, our results still hold in the extended setting of Eq. (2).

2.2. Maximization problem. In this subsection we present a stylized game theoretical model that leads to our equation of interest, Eq. (3). It is another justification of Eq. (3) along with the primary one, which is to capture a number of essential properties of networks (non-negativity of edges, positive probability of vanishing of each edge, and interactions between edges, which affect the whole network).

Consider a world with n myopic agents (people/firms/countries/etc.) with quadratic adjustment costs. Agents can interact with each other over time. Time is discrete and goes from 1 to T . Every period, each agent i chooses how much time to spend with or how much to trade with each other agent j . The decision is based on two components: costs and benefits.

Benefits are characterized by a per unit gain of $\alpha_{ij} + u_{ijt}$. Here α_{ij} is a constant, while u_{ijt} is random component that is independent across time. Thus, y units of communication/trade leads to a benefit of $y(\alpha_{ij} + u_{ijt})$.

The second component is a quadratic adjustment cost function. Agents get disutility whenever there are deviations from some target expected level of communication/trade. The target is composed from an own past and a peer-effect or interactions component. The interaction term aggregates the whole structure of the network for up to H periods. That is, we assume that agent i by choosing to devote y units to agent j has to pay

$$\frac{1}{2} \left(y - \beta_{ij}y_{ijt-1} - \gamma_{ij}p_{ij} \left(\{y_{kls}\}_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} \right) \right)^2,$$

where

$$p_{ij} \left(\{y_{kls}\}_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} \right) : \mathbb{R}_+^{n^2H} \rightarrow \mathbb{R}$$

represents peer effects/interactions function, which depends on H previous periods.

The interpretation is that β_{ij} represents a rate at which stock/relationship depreciates/appreciates. If $\beta_{ij} < 1$, then the agent is introverted and tends to decrease communication; while if $\beta_{ij} > 1$, the agent is an extrovert, who tends to expand communication. The interpretation of β_{ij} for firms corresponds to production depreciation (i.g. technology wears out) or production appreciation (better technology management over time increases production). The coefficient γ_{ij} indexes the sensitivity of the reference level with respect to peer effects/interactions, which are, in turn, represented by the function p_{ij} . The peer effects function depends on H past periods of the whole network, and captures interactions between different edges y_{kls} across time.

Agents have separate maximization problems for each time period t and with each peer j . Agent i solves the following maximization problem at day t with respect to agent j :

$$(4) \quad \max_{y \geq 0} \left[y(\alpha_{ij} + u_{ijt}) - \frac{1}{2} \left(y - \beta_{ij}y_{ijt-1} - \gamma_{ij}p_{ij} \left(\{y_{kls}\}_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} \right) \right)^2 \right].$$

The solution to the maximization problem (4) is

$$y_{ijt}^* = \left[\alpha_{ij} + \beta_{ij}y_{ijt-1} + \gamma_{ij}p_{ij} \left(\{y_{kls}\}_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} \right) + u_{ijt} \right]_+,$$

which leads to the network evolution process described by Eq. (3).

2.3. Assumptions. We need to impose some assumptions on the error distribution and on the peer effects function.

Assumption 1. *The vector $\{u_{ijt}\}_{i,j=1,\dots,n}$ is i.i.d. over t .*

Assumption 1 allows errors to have arbitrary correlation across edges of the network and ensures that those correlations do not vary over time. This differs substantially from classical approach to networks, where they are modeled as cross-section or short panel. In those cases the correlation across individuals must either form a dyadic relationship or decrease with some measure of a distance, so that one can average over individuals and apply the law of large numbers. In our setting n is assumed to be fixed, and instead the variation across T is used to identify the network formation process. Thus, our setting does not rely on “decreasing across n correlations” and can accommodate any type of cross-sectional interdependence.

Assumption 2. *The support of $\{u_{ijt}\}_{i,j=1,\dots,n}$ satisfies $\mathbb{P}(u_{ijt} < -M \forall i, j) > 0$ for all $M > 0$.*

Assumption 2 is used to show stationarity. It implies that errors jointly take large negative values with positive probability. The consequence is that for any values of the process $\{y_{ijt}\}_{i,j}$ at time t , there is a positive probability that at time $t + 1$ the process jumps to zero. This observation is crucial for Theorems 2 and 3, as it ensures that for specific range of parameters the process forgets its past in a finite time and restarts from scratch.

Remark 2. Convergence to stationarity will be used to show consistency and asymptotic normality of our estimators of $\{\alpha_{ij}, \beta_{ij}, \gamma_{ij}\}_{i,j}$. Yet, it is plausible that one can relax Assumption 2, and still get a consistent estimator as long as several first moments of the network evolution process $\{y_{ijt} : i, j = 1 \dots, n\}_{t \geq 1}$ are uniformly bounded. Theorem 1 guarantees uniformly bounded first moment and analogues results can be obtained for higher order moments. We further discuss estimation without stationarity in Remark 12.

Eq. (3) has one degree of freedom, so we need to impose a normalization assumption on the error term u_{ijt} . We consider two different normalization assumptions: zero mean or zero median, which are stated below.

- (1) *Normalization of the mean:* For all i, j, t , $\mathbb{E}u_{ijt} = 0$.
- (2) *Normalization of the median:* For all i, j, t , $\text{med}(u_{ijt}) = 0$.

Alternative normalizations only lead to differences in α_{ij} :

$$\alpha_{ij}^E = \alpha_{ij}^{\text{med}} + \mathbb{E}u_{ijt}^{\text{med}},$$

where α_{ij}^E is an intercept under normalization of the mean and α_{ij}^{med} and $\mathbb{E}u_{ijt}^{\text{med}}$ are an intercept and a mean of the error under normalization of the median.

Assumption 3 (Peer effects do not grow faster than their maximal argument.).

$p : \mathbb{R}_+^{n^2 H} \rightarrow \mathbb{R}$ is such that there exists a constant $\mathcal{A} \in \mathbb{R}$ for which

$$\left| p \left(\left\{ y_{kls} \right\}_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} \right) \right| \leq \mathcal{A} + \max_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} y_{kls}.$$

Remark 3. The choice of constant 1 in front of $\max_{k,l,s} y_{kls}$ is without loss of generality. Writing instead $C \max_{k,l,s} y_{kls}$, where $C > 0$ is some constant, only leads to rescaling of the peer-effect function and its coefficient γ_{ij} .

Let us present some examples of possible peer effect functions.

- *Maximum:*

$$p_{ij} \left(\left\{ y_{kls} \right\}_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} \right) = \max_{\substack{(k,l) \neq (i,j) \\ s=t-H,\dots,t-1}} y_{kls}.$$

This function represents the largest possible stimulus to increase trade or communication. This can be interpreted as a steadily expanding economy.

- *Minimum:*

$$p_{ij} \left(\left\{ y_{kls} \right\}_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} \right) = \min_{\substack{(k,l) \neq (i,j) \\ s=t-H,\dots,t-1}} y_{kls}.$$

This function corresponds to the smallest, but still non-zero influence from others. That is, if some edge jumps to zero, it pushes the other edges in that direction. Alternatively, if all edges have positive weights, the peer effect term is still positive and helps to maintain a non-zero edge between i and j .

- *Linear:*

$$p_{ij} \left(\left\{ y_{kls} \right\}_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} \right) = \sum_{\substack{k,l=1,\dots,n \\ r=1,\dots,H}} \lambda_{klr} y_{klt-r},$$

$$\text{where } \{\lambda_{klr}\}_{\substack{k,l=1,\dots,n \\ r=1,\dots,H}} \text{ are known and } \sum_{\substack{k,l=1,\dots,n \\ r=1,\dots,H}} |\lambda_{klr}| \leq 1.$$

The linear function represents an intermediate point between the previous examples.

- *Triangles:*

$$p_{ij} \left(\left\{ y_{kls} \right\}_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} \right) = \sum_{k \neq i,j} \frac{\sqrt{y_{ikt-H} y_{kjt-H}}}{n-2}.$$

The square root makes the order of the peer effect term to be the same as the order of the autoregressive term, y_{ijt} . Thus, Assumption 3 is satisfied.

Triangular peer effects have the most interesting functional form, and we use it in our empirical application. The interpretation is that if i and k are strongly connected, and k and j are too, then there is a higher probability of a connection between i and

j in the future. Note that the product makes it important that both connections are present. If k and j are not connected, we cannot expect k to “introduce j to i ”. Thus, we look at all triangles which have (i, j) as one of their legs. Two strong links in such triangles are expected to strengthen the third leg, (i, j) . This can be summed up as “friend of my friend is my friend”. Such peer effects may be present in social interactions, interactions between firms or countries, and so on.

The triangular peer effect function is related to the network formation process in Graham (2016), where the presence of an edge between two nodes depends on the number of triangles containing those two nodes in the past. In Graham (2016) edges do not have weights, so the number of triangles containing i and j at time t is $\sum_k D_{ikt} D_{kjt}$, where $D_{ikt} = 1$ when there is an edge between i and k at time t . A similar statistics $\max_k D_{ikt} D_{kjt}$ is used in Leung and Moon (2019) to generate network clustering.

2.4. Stationarity. Theorem 1, which is proved in the Supplementary Material, provides sufficient conditions under which the network does not explode. Non-exploding does not guarantee convergence to a stationary distribution as formally the process may have cycles. Yet, it is enough for our estimation and prediction approaches to work. Moreover, if we additionally assume that $\{u_{ijt}\}_{i,j}$ has unbounded from below support (Assumption 2), then stationarity holds (Theorem 2).

Theorem 1. *Suppose that Assumptions 1 and 3 are satisfied, $\mathbb{E}u_{ijt}$ exists for all i, j, t , and there exists a constant $C \in (0, 1)$ such that $\max(0, \beta_{ij}) + |\gamma_{ij}| < C$ for all i, j . Then the multivariate process $\{y_{ijt} : i, j = 1 \dots, n\}_{t \geq 1}$ does not explode (i.e. there exists a constant C_1 such that $\mathbb{E}y_{ijt} < C_1 < \infty$ for all i, j, t).*

Definition. *The process $\bar{y}_t = \{y_{ijt}\}_{i,j}$ is strongly mixing if for arbitrary Borel sets Δ_1, Δ_2*

$$\lim_{t \rightarrow \infty} |\mathbb{P}(\bar{y}_s \in \Delta_1, \bar{y}_{t+s} \in \Delta_2) - \mathbb{P}(\bar{y}_s \in \Delta_1)\mathbb{P}(\bar{y}_{t+s} \in \Delta_2)| = 0.$$

Theorem 2. *Suppose that Assumptions 1, 2, and 3 are satisfied, $\mathbb{E}u_{ijt}^4 < \infty$ for all i, j, t , and there exists a constant $C \in (0, 1)$ such that $\max(0, \beta_{ij}) + |\gamma_{ij}| < C$ for all i, j , then the*

multivariate process $\{y_{ijt} : i, j = 1 \dots, n\}_{t \geq 1}$ is strongly mixing and converges to a stationary process.²

Remark 4. In the extended setting of Eq. (2), Assumption 3 should hold for each p_{ij}^κ , and $|\gamma_{ij}|$ in Theorem 2 should be replaced by the sum $\sum_{\kappa=1}^K |\gamma_{ij}^\kappa|$.

A striking feature of Theorem 2 is that we do not need the error distribution to be continuous to get strong mixing. This differs from the linear case (see Withers (1981) and Andrews (1984) for examples of $AR(1)$ processes which are not strongly mixing). The reason is that in our setting the expected time until the process jumps to be identically zero ($y_{ijt} = 0$ for all i, j) is finite. Thus, the process forgets the initial condition in finite time.

Example 1. The fact that the peer effects function p_{ij} can depend only on a fixed number of time periods is crucial. For example, suppose that we have only one equation ($n = 2$) which is initialized at $y_0 = 0$, and the error process u_t has unbounded support from above. Further suppose $\alpha = \beta = 0$, $\gamma = 0.5$ and $z_t := p_{ij}(y_t, \dots, y_0) = \max(y_t, \dots, y_0)$, so that Assumption 3 is satisfied. Then $\beta + |\gamma| = 0.5 < 1$, but the process y_t is explosive. To see this, let us analyze the behavior of y_t and z_t .

By definition, $z_0 = y_0 = 0$, so that $y_1 = [u_1]_+$ and $z_1 = \max(0, u_1) \geq 0$. Thus, $y_2 = [0.5z_1 + u_2]_+ \geq [u_2]_+$ and $z_2 = \max(0, y_1, y_2) \geq \max(0, u_1, u_2)$. Similarly, $y_3 = [0.5z_2 + u_3]_+ \geq [u_3]_+$ so that $z_3 \geq \max(0, u_1, u_2, u_3)$. Applying induction, for any t we get $z_t \geq \max(0, u_1, \dots, u_t)$. Therefore, $z_t \xrightarrow[t \rightarrow \infty]{a.s.} \infty$, as the support of u_t is unbounded from above and the maximum of an infinite number of random variables with unbounded support diverges. Because $y_t = [0.5z_{t-1} + u_t]_+$, y_t also goes to infinity almost surely.

3. Special case: No interactions

In this section we consider a special case, where the connection between i and j at time t depends only on its past. That is, past interactions between $k \neq i, j$ and l do not influence

²Formally this means that the finite-dimensional distributions of the process $\{y_{ijt+\tau} : i, j = 1 \dots, n\}_{\tau \in \mathbb{Z}}$, converge to those of a stationary in τ process as $t \rightarrow \infty$.

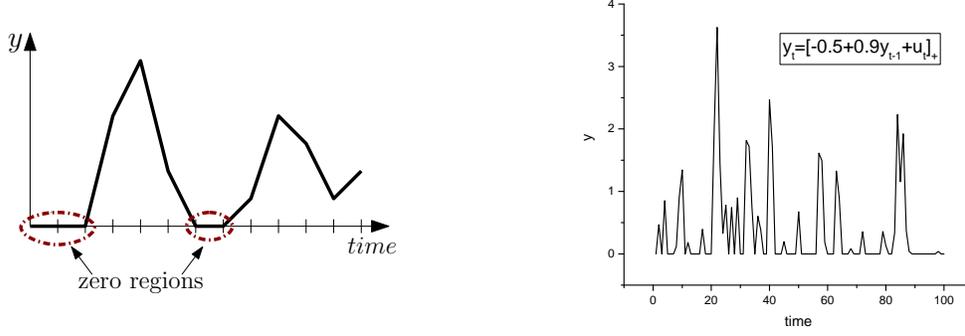


FIGURE 1. A typical sample path for $y_t = [\alpha + \beta y_{t-1} + u_t]_+$, $u_t \sim i.i.d.(0, \sigma^2)$.

y_{ijt} . Thus, the model reduces to $\frac{n(n-1)}{2}$ separate equations of the form

$$(5) \quad y_t = [\alpha + \beta y_{t-1} + u_t]_+, \quad u_t \sim i.i.d.(0, \sigma^2).$$

A typical sample path for Eq. (5) is shown in Figure 1. When the process hits zero, it stays at zero for some time, then goes to an “AR(1)-excursion”, until it becomes negative. Positive part in (5) then forces y_t to become zero instead, and everything starts again.

The following theorem provides a full classification of stationary/explosive behavior in the case of no interactions. In contrast to classical autoregression which has no positive part, when $\beta = 1$ the process still converges to a stationary distribution when $\alpha < 0$.

Theorem 3. (*Classification Theorem*) *Let Assumptions 1 and 2 hold. Assume further that $\mathbb{P}(u_t > -\alpha) > 0$, $\mathbb{E}u_t = 0$, and $\mathbb{E}u_t^4 < \infty$,*

- *If $\beta < 1$, then y_t is strongly mixing and converges to a stationary process;*
- *If $\beta = 1$, $\alpha < 0$, then y_t is strongly mixing and converges to a stationary process;*
- *If $\beta > 1$, then y_t is divergent: y_t converges to ∞ almost surely;*
- *If $\beta = 1$, $\alpha > 0$, then y_t is divergent: y_t converges to ∞ almost surely;*
- *If $\beta = 1$, $\alpha = 0$, then y_t is mean-divergent: $\mathbb{E}y_t \rightarrow \infty$. The proper scaling limit is*

$$\frac{1}{\sqrt{T}}y_{\lfloor Tr \rfloor} \xrightarrow{d} \sigma|W(r)|, \quad r \in [0, 1], \quad \text{as } T \rightarrow \infty,$$

where $W(\cdot)$ is a standard Brownian motion and $\mathbb{E}u_t^2 = \sigma^2$.

A visual summary of the results in Theorem 3 is shown in Figure 2, where the evolution of y_t is illustrated for different values of α and β .

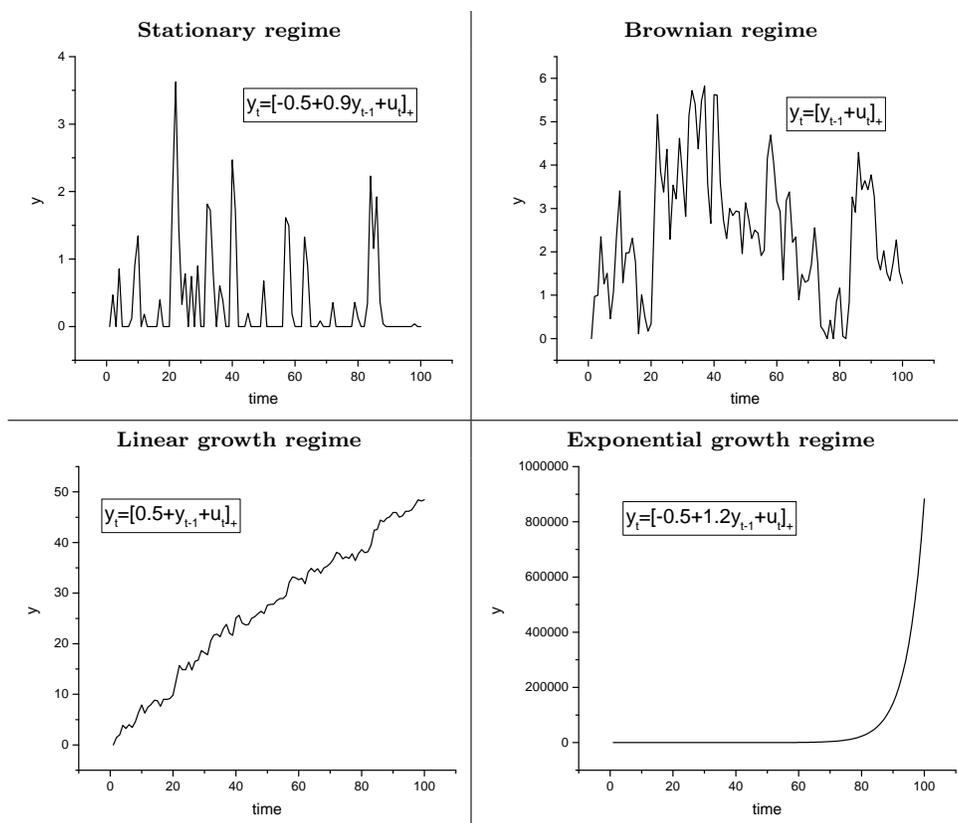


FIGURE 2. Illustration of Theorem 3.

Theorem 3 is proved in the Section C of the Supplementary Material (Theorems C.4, C.5, C.6, and C.8). The stationarity part of the proof relies on the large deviations principle and the renewal theorem. The idea is to show that the expected time until the process reaches zero is finite. Then one can apply the renewal theorem to get the limiting distribution. Interestingly, just the knowledge that the process hits zero with probability one is not enough. In particular, if $\alpha = 0, \beta = 1$ we get a standard unit-root process for time periods when y_t is positive. This process always hits zero. However, it does not converge to a stationary distribution, and as Theorem 3 shows, the process $y_t = [y_{t-1} + u_t]_+$ has exploding mean. Thus, there is a discontinuity between the stationary and explosive regions. This is similar to what is observed in the classical linear autoregressive case. However, while in the classical case the re-scaled unit root process converges to the Brownian motion, in our truncated setting we get the absolute value of the Brownian motion, see Theorem C.8 in the Supplementary Material and Cavaliere (2004).

The limiting distribution in the stationary case of Theorem 3 is complicated and cannot be written explicitly as a function of α , β , and the distribution of u . However, what can be calculated is the expected time the process y_t spends at zero once it hits zero. In particular, the average length of an interval of zeros is decreasing in α .

Lemma 4. *Let Assumption 1 hold. Once the process y_t hits zero, the expected time it spends until finally jumping to a positive value is $\frac{1}{1-F_u(-\alpha)}$, where F_u is the cdf of u_t .*

Proof. If $y_t = 0$, then $y_{t+1} > 0$ when $u_{t+1} > -\alpha$. Therefore, after each zero with probability $F_u(-\alpha)$ the process remains at zero and with the remaining probability it becomes positive. Thus, we get a sequence of Bernoulli random variables, where the expected time until the first tail observation is

$$1 \cdot (1 - F_u(-\alpha)) + 2 \cdot F_u(-\alpha)(1 - F_u(-\alpha)) + 3 \cdot F_u^2(-\alpha)(1 - F_u(-\alpha)) + \dots = \frac{1}{1 - F_u(-\alpha)}. \quad \square$$

4. Estimation

We now return back to the general model summarized in Eq. (3). The two main difficulties in our model are the interactions between the outcome variables and the non-separable errors. We can overcome the former by estimating the following model

$$(6) \quad y_t = [\alpha + \beta y_{t-1} + \gamma z_{t-1} + u_t]_+.$$

To be more specific, for each pair (i, j) we define

$$z_{ijt-1} = p_{ij} \left(\left\{ y_{kls} \right\}_{\substack{k,l=1,\dots,n \\ s=t-H,\dots,t-1}} \right)$$

and ignore the fact that z_{ijt} are functions of lags of y_{klt} , $k, l = 1, \dots, n$. That is, after we calculate the values of z_{ijt} , we are not going to use the fact that those values were obtained from $\{y_{klt}\}_{k,l=1,\dots,n}$ and its lages. Instead we treat z_{ijt} as any other regressors. For each edge $(i \rightarrow j)$ we separately estimate Eq. (6) with $y_t = y_{ijt}$, $z_t = z_{ijt}$.

The natural question is whether we lose predictive power by treating each equation independently or not. Generally it is not clear and the answer should depend on the class of estimators we consider. Yet, one could hope that we do not lose a lot.³

In the following subsections we present an approach to estimating the model (6) and discuss the properties of the estimator.

Remark 5. All of the results in this section can be straightforwardly generalized to the case of multiple peer effects terms. I.e., to the model with K regressors $z_{t-1}^1, \dots, z_{t-1}^K$

$$y_t = [\alpha + \beta y_{t-1} + \sum_{k=1}^K \gamma^k z_{t-1}^k + u_t]_+,$$

where coefficients $\alpha, \beta, \{\gamma^k\}_{k=1}^K$ are unknown and have to be estimated. For example, for the above model, the analogue of the matrix M_R defined in Eq. (8) and used in Theorems 5 and 6 is a $K + 2$ by $K + 2$ matrix composed of all second moments of the vector $(1, y_t, z_t^1, \dots, z_t^K) \mathbf{1}(\alpha + \beta y_t + \sum_{k=1}^K \gamma^k z_t^k \geq R)$.

4.1. L₁ Estimation. The least absolute deviations (LAD) estimator is the solution to the following minimization problem

$$(7) \quad \min_{a,b,c} \sum_{t=1}^T |y_t - [a + by_{t-1} + cz_{t-1}]_+|.$$

It turns out, as Theorems 5 and 6 show, the LAD estimators are consistent and asymptotically normal. In what follows (y, z) denotes the distributional limit of (y_t, z_t) as $t \rightarrow \infty$.

Let $R \geq 0$ and define the matrix M_R as

$$(8) \quad M_R = \mathbb{E} \left[\begin{pmatrix} 1 & y & z \\ y & y^2 & yz \\ z & yz & z^2 \end{pmatrix} \mathbf{1}(\alpha + \beta y + \gamma z \geq R) \right],$$

³There are examples of such statements in the literature. E.g., Gouriou and Monfort (1980) show that for the linear models like vector autoregressions if the errors are independent across components of y_t (in our case across edges) or the matrices of regressors span the same subspace for each component of y_t , then single-equation generalized least squares (GLS) is equivalent to overall GLS.

Theorem 5. *Assume that y_t is strongly mixing and converges to a stationary distribution and $(\alpha, \beta, \gamma) \in \Theta$, where Θ is some compact set in \mathbb{R}^3 . If u_t is i.i.d. with a continuous and positive density at 0, $\text{med}(u_t) = 0$, and M_R is nonsingular for some $R > 0$ at the true parameter values, then the LAD estimator is consistent:*

$$\left(\hat{\alpha}_{LAD}, \hat{\beta}_{LAD}, \hat{\gamma}_{LAD} \right) \xrightarrow[T \rightarrow \infty]{\mathbb{P}} (\alpha, \beta, \gamma).$$

Remark 6. The singularity of the matrix M_R is a rare event. The matrix M_R is singular if $\beta = \gamma = 0$, $\alpha < R$. In this case the indicator equals zero, so the matrix is identically zero. The other case is if $\beta = 0$, $\gamma > 0$ and z never takes values above $\frac{R-\alpha}{\gamma}$, so that again the matrix M_R is identically zero.

If $\beta > 0$, the peer effect function $p_{ij}(\cdot)$ does not depend on y_{ijt} , and the random variable z is non-constant, then the matrix is nonsingular. For instance, the triangular peer effect functions p_{ij} do not depend on y_{ijt} . Minimum and maximum functions, if taken over all edges except the given edge ($i \rightarrow j$), also do not depend on y_{ijt} . Similarly linear functions satisfy this as long as the corresponding weight $\lambda_{ijt} = 0$. Thus, in all those examples M_R is non-singular for $\beta > 0$.

The idea of the above argument is that if M_R is singular, then there exists a non-zero vector $(\lambda_1, \lambda_2, \lambda_3)$ such that $\lambda_1 \mathbf{1}(\alpha + \beta y + \gamma z \geq R) + \lambda_2 y \mathbf{1}(\alpha + \beta y + \gamma z \geq R) + \lambda_3 z \mathbf{1}(\alpha + \beta y + \gamma z \geq R) \equiv 0$. That is, when $\alpha + \beta y + \gamma z \geq R$, we must have $\lambda_1 + \lambda_2 y + \lambda_3 z = 0$. However, as $\beta > 0$ and z does not depend on y , we can perturb y a bit and get $y' = y + \varepsilon$, $\varepsilon > 0$, in which case the indicator is still non-zero, but the second equality fails unless $\lambda_2 = 0$. If $\lambda_2 = 0$, then we must have $z = -\lambda_1/\lambda_3$ whenever the indicator equals one. This again is impossible, when z is not a fixed constant.

Remark 7. Note that if $R_1 > R_2$ and M_{R_1} is nonsingular, then M_{R_2} is also nonsingular. Equivalently, if M_{R_2} is singular, then so is M_{R_1} . The reason is that if there exists a non-zero vector $(\lambda_1, \lambda_2, \lambda_3)$ such that $\lambda_1 + \lambda_2 y + \lambda_3 z = 0$ when $\alpha + \beta y + \gamma z \geq R_2$, then the same holds for $\alpha + \beta y + \gamma z \geq R_1$, as $R_1 > R_2$. Thus, there is a bound $\bar{R} \in \mathbb{R} \cup \{+\infty\}$ such that M_R is nonsingular for any $R < \bar{R}$ and singular for any $R > \bar{R}$.

Remark 8. The condition on the error density $f_u(0) > 0$ can be weakened to

$$\mathbb{P}(u \in [-\varepsilon, 0)) > 0, \mathbb{P}(u \in (0, \varepsilon]) > 0 \text{ for any } \varepsilon > 0.$$

That is, u_t must be in the left and right neighbourhoods of zero with positive probability. From the proof of Theorem 5, we only need $\min \left(\int_{-\tau}^0 (\tau + u) dF_u(u), \int_0^{\tau} (\tau - u) dF_u(u) \right)$ to be positive for any $\tau > 0$, which is satisfied in this case.

On the other hand, for Theorems 6 and 8 such weakening of the condition $f_u(0) > 0$ leads to a change in the asymptotic theory, and we do not address it in the present paper.

Theorem 6. *Assume that y_t is strongly mixing and converges to a stationary distribution and $(\alpha, \beta, \gamma) \in \Theta$, where Θ is some compact set in \mathbb{R}^3 . If u_t is i.i.d. with a continuous and positive density at 0, $\text{med}(u_t) = 0$, M_0 is nonsingular at the true parameter values, and the random function $\Delta \mapsto \mathbf{1}(\alpha + \beta y + \gamma z + \Delta > 0)$ is continuous with probability one at the true parameter values and at $\Delta = 0$, then the LAD estimator is asymptotically normal:*

$$\sqrt{T} \begin{pmatrix} \hat{\alpha}_{LAD} - \alpha \\ \hat{\beta}_{LAD} - \beta \\ \hat{\gamma}_{LAD} - \gamma \end{pmatrix} \xrightarrow[T \rightarrow \infty]{d} \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \frac{1}{4f_u^2(0)} M_0^{-1} \right).$$

Remark 9. The condition that the random function $\Delta \mapsto \mathbf{1}(\alpha + \beta y + \gamma z + \Delta > 0)$ is continuous with probability one at the true parameter values and at $\Delta = 0$ is generally hard to violate, since the violation requires $\alpha + \beta y + \gamma z$ to equal zero with positive probability. This can happen when $\alpha = 0$, because $y = 0$ with positive probability. If z is also zero with positive probability, then the indicator jumps depending on the sign of Δ . On the other hand if $\alpha > 0$, $\beta, \gamma \geq 0$ and $z > 0$ with probability 1, then the indicator is identical 1. This corresponds to the special case of Theorem 7.

Remark 10. We do not impose a condition on the existence of nonsingular M_R for some $R > 0$ in Theorem 6. This is because non-degeneracy of M_0 together with continuity of $\mathbf{1}(\alpha + \beta y + \gamma z + \Delta > 0)$ imply the existence of $R > 0$, for which M_R is also nonsingular.

In the special “non-negative” case the conditions of Theorems 5 and 6 can be simplified: $[a + by_{t-1} + cz_{t-1}]_+ \equiv a + by_{t-1} + cz_{t-1}$ when a is positive and b, c and z_{t-1} are non-negative (y_{t-1} is non-negative by assumption). Thus, in this case the positive part in the minimization problem never binds, and the additional condition on the stationary distribution and true parameter does not arise.

Moreover, the minimization problem (7) is convex when $a > 0, b \geq 0, c \geq 0$, and $z \geq 0$. Thus, the numerical solution is a global maximum when the true $\alpha > 0, \beta \geq 0, \gamma \geq 0$.

Theorem 7. *Assume that y_t is strongly mixing and converges to a stationary distribution. If $\alpha > 0, \beta \geq 0, \gamma \geq 0$, the peer effect function is non-negative (i.e., $z_t \geq 0$ for all t), u_t is i.i.d. with a continuous and positive density at 0, $\text{med}(u_t) = 0$, and the random variables $1, y, z$ are linearly independent⁴, then the LAD estimator is consistent:*

$$\left(\hat{\alpha}_{LAD}, \hat{\beta}_{LAD}, \hat{\gamma}_{LAD} \right) \xrightarrow{\mathbb{P}} (\alpha, \beta, \gamma) \text{ as } T \rightarrow \infty.$$

Remark 11. The linear independence condition holds automatically whenever the peer effect function $p_{ij}(\cdot)$ depends on an argument other than y_{ijt-1} in a non-degenerate way. Indeed, in this case y_t contains only error u_{ijt} , while z_t also contains errors u_{klt} for $(k, l) \neq (i, j)$ (if z_t only depends on lagged peer effects, then it is independent of u_{ijt}). Thus, there is uncertainty, which cannot be removed by taking linear combinations.

Theorem 8. *Assume that y_t is strongly mixing and converges to a stationary distribution. If $\alpha > 0, \beta \geq 0, \gamma \geq 0$, the peer effect function is non-negative (i.e., $z_t \geq 0$ for all t), u_t is i.i.d. with a continuous and positive density at 0, $\text{med}(u_t) = 0$, and the random variables $1, y, z$ are linearly independent, then the LAD estimator is asymptotically normal:*

$$\sqrt{T} \begin{pmatrix} \hat{\alpha}_{LAD} - \alpha \\ \hat{\beta}_{LAD} - \beta \\ \hat{\gamma}_{LAD} - \gamma \end{pmatrix} \xrightarrow[T \rightarrow \infty]{d} \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \frac{1}{4f_u^2(0)} \begin{pmatrix} 1 & \mathbb{E}y & \mathbb{E}z \\ \mathbb{E}y & \mathbb{E}y^2 & \mathbb{E}yz \\ \mathbb{E}z & \mathbb{E}yz & \mathbb{E}z^2 \end{pmatrix}^{-1} \right).$$

⁴Linear independence means that there do not exist constants $\lambda_1, \lambda_2, \lambda_3$ such that $\lambda_1 + \lambda_2 y + \lambda_3 z = 0$ almost surely (1 represents a constant random variable).

Remark 12. Theorems 5–8 assume that y_t is strongly mixing and converges to a stationary distribution. This is guaranteed by Theorems 2 and 3 under Assumption 2 (unbounded from below support of errors). We conjecture that if y_t does not converge to a stationary distribution, but has uniformly bounded several first moments⁵, consistency and asymptotic normality still hold, as long as we replace M_0 with its finite sample version

$$\widehat{M}_0 = \frac{1}{T} \sum_t \begin{pmatrix} 1 & y_t & z_t \\ y_t & y_t^2 & y_t z_t \\ z_t & y_t z_t & z_t^2 \end{pmatrix} \mathbf{1}(\alpha + \beta y_t + \gamma z_t \geq 0).$$

That is, we conjecture that

$$\sqrt{T} \widehat{M}_0^{1/2} \left(\hat{\alpha}_{LAD} - \alpha, \hat{\beta}_{LAD} - \beta, \hat{\gamma}_{LAD} - \gamma \right)' \xrightarrow[T \rightarrow \infty]{d} \frac{1}{2f_u(0)} \mathcal{N}(0, I).$$

4.2. LAD in the explosive case. For the model without peer effects Theorem 3 provides a full classification of the asymptotic behavior of y_t . Using this theorem it is possible to establish consistency of the LAD estimator in the model without peer effects not only under a stationary distribution, but also for explosive and mean-explosive scenarios. This corresponds to cases when $\beta > 1$ or $\{\beta = 1, \alpha + \mathbb{E}u_t \geq 0\}$.⁶

Theorem 9. *Suppose $\gamma = 0$ (no peer effects) and suppose that $(\alpha, \beta) \in \Theta$, where Θ is some compact set in \mathbb{R}^2 . If u_t is i.i.d. with a continuous and positive density at 0, $\text{med}(u_t) = 0$, then the LAD estimator is consistent for $\{\beta = 1, \alpha + \mathbb{E}u_t \geq 0\}$ and $\beta > 1$:*

$$\left(\hat{\alpha}_{LAD}, \hat{\beta}_{LAD}, \right) \xrightarrow[T \rightarrow \infty]{\mathbb{P}} (\alpha, \beta).$$

We conjecture that consistency holds for any value of γ . Because the proof of Theorem 9 uses a separate argument depending on the asymptotic behavior of y_t , we cannot extend it to the model with peer effects. To be more precise, when $\gamma = 0$ the model has three types of explosive behavior: exponential growth regime ($\beta > 1$), linear growth regime ($\beta =$

⁵E.g., see Theorem 1 which guarantees uniformly bounded first moment.

⁶Notice that in Theorem 3 we had α instead of $\alpha + \mathbb{E}u_t$. This is due to the fact that we assumed $\mathbb{E}u_t = 0$ in Theorem 3, while now we use a different normalization, $\text{med}(u_t) = 0$.

$1, \alpha + \mathbb{E}u_t > 0$), and Brownian regime ($\beta = 1, \alpha + \mathbb{E}u_t = 0$). Yet, it is unclear what the analogue of these regimes is when $\gamma \neq 0$.

4.3. Discussion: LAD vs OLS. Interestingly, although the model is non-linear we still can treat it as linear to get consistent L_1 or LAD estimates. Yet, the same approach does not work with L_2 or OLS. The following example illustrates that OLS leads to inconsistency bias, while LAD does not.

Example 2. Let $y_t = [\alpha + u_t]_+$, $a \geq 0$, $\text{med}(u_t) = \mathbb{E}u_t = 0$. Then the LAD estimate solves

$$\min_{\hat{\alpha}} \sum_t |y_t - \hat{\alpha}|.$$

The solution to the minimization problem is the sample median, that is $\hat{\alpha}_{LAD} = \text{med}(y_1, \dots, y_T)$. As $T \rightarrow \infty$ the sample median converges to the median of the stationary distribution of y_t . The median of y_t equals α , because with probability 0.5 the error u_t is positive, so that $y_t = \alpha + u_t \geq \alpha$, and with probability 0.5 the error u_t is negative, so that either $y_t = 0 \leq \alpha$ or $y_t = \alpha + u_t \leq \alpha$. Thus, with probability 0.5 $y_t \geq \alpha$ and with probability 0.5 $y_t \leq \alpha$. So $\hat{\alpha}_{LAD} \xrightarrow[T \rightarrow \infty]{\mathbb{P}} \text{med}(y) = \alpha$.

However, the results are different if we minimize an L_2 norm instead of an L_1 norm ignoring positive part. The solution to

$$\min_{\hat{\alpha}} \sum_t (y_t - \hat{\alpha})^2$$

is the sample mean, $\hat{\alpha}_{OLS} = \frac{\sum_t y_t}{T}$. As $T \rightarrow \infty$ the sample mean converges to the actual mean of y_t , so that

$$\hat{\alpha}_{OLS} \xrightarrow[T \rightarrow \infty]{\mathbb{P}} \mathbb{E}y_t = \mathbb{E}[\alpha + u_t]_+ = \int_{-\alpha}^{\infty} (\alpha + u) f_u(u) du = \alpha(1 - F(-\alpha)) + \int_{-\alpha}^{\infty} u f_u(u) du \neq \alpha.$$

The intuition is that median is more robust to truncation at zero: if the median of a process is positive, it does not matter if we replace negative values with zero and vice versa. Yet, the mean is significantly shifted by such procedure.

In the Supplementary Material we show how to correct the OLS procedure to get consistent estimates. The idea is to use “identification from infinity” approach, cf. Chamberlain (1986).

5. Prediction

In this section we again treat equations for each edge separately. After the parameters of the model are estimated, one can do predictions. The model is Markovian (in a state space taking into account H time lags), thus, can be easily used for predictions. This leads us to propose a prediction of the form $\hat{y}_{T+1}^{LAD} = [\hat{\alpha}_{LAD} + \hat{\beta}_{LAD}y_T + \hat{\gamma}_{LAD}z_T]_+$.

We use a rolling window approach to measure predictive power: we choose some number $T' < T$, and estimate the model based on observations $t, \dots, t + T' - 1$. Then we form a forecast $\hat{y}_{t+T'}$ based on the obtained estimates and observed values $\{y_t, \dots, y_{t+T'-1}\}$. For each $t = 1, \dots, T - T'$ we calculate the difference between the predicted out-of-sample value $\hat{y}_{t+T'}$ and the actual value $y_{t+T'}$. We get a measure of how well we can predict the data:

$$(9) \quad R_{abs} = \frac{1}{T - T'} \sum_{t=1}^{T-T'} |y_{t+T'} - \hat{y}_{t+T'}|.$$

The smaller R_{abs} is, the better predictions we have on average. Similarly, we can also sum over all pairs (i, j) to get a prediction measure over the whole network.

There are two reasons why we use absolute deviations, i.e. the L_1 norm and not the more usual L_2 norm. First, as the estimation relies on minimizing the L_1 norm, it is more consistent to also use the same norm to evaluate predictions. Second, the 1-step-ahead optimal prediction in L_2 norm when α , β , and γ are known is $\int_{-\alpha - \beta y_{t-1} - \gamma z_{t-1}}^{\infty} (\alpha + \beta y_{t-1} + \gamma z_{t-1} + u) f_u(u) du$. Thus, it crucially depends on the distribution of the error term, f_u , even when there is no parameter uncertainty. However, the 1-step-ahead optimal prediction in L_1 when α , β , and γ are known is $[\alpha + \beta y_{t-1} + \gamma z_{t-1}]_+$, as shown below. That is, it does not depend on the distribution of the error and is more convenient to work with.

Remark 13. Suppose that α , β , and γ are known. To see that the optimal 1-step-ahead prediction in L_1 is $[\alpha + \beta y_{t-1} + \gamma z_{t-1}]_+$, define $\Delta \hat{y}_{t+1} = \hat{y}_{t+1} - \alpha - \beta y_t - \gamma z_t$ and write

$$(10) \quad \begin{aligned} \int |y_{t+1} - \hat{y}_{t+1}| f(u) du &= \int |[\alpha + \beta y_t + \gamma z_t + u]_+ - \hat{y}_{t+1}| f(u) du \\ &= \int |\max(u, -\alpha - \beta y_t - \gamma z_t) - \Delta \hat{y}_{t+1}| f(u) du. \end{aligned}$$

Because $\arg \min_C \mathbb{E}|v - C| = \text{med}(v)$, and

$$\max(u, -\alpha - \beta y_t - \gamma z_t) = \begin{cases} u, & u \geq -\alpha - \beta y_t - \gamma z_t, \\ -\alpha - \beta y_t - \gamma z_t, & u < -\alpha - \beta y_t - \gamma z_t, \end{cases},$$

$$\text{med}(\max(u, -\alpha - \beta y_t - \gamma z_t)) = \begin{cases} 0, & \alpha + \beta y_t + \gamma z_t \geq 0, \\ -\alpha - \beta y_t - \gamma z_t, & \alpha + \beta y_t + \gamma z_t < 0, \end{cases}$$

minimizing Eq. (10), gives

$$\Delta \hat{y}_{t+1} = \begin{cases} 0, & \alpha + \beta y_t + \gamma z_t \geq 0, \\ -\alpha - \beta y_t - \gamma z_t, & \alpha + \beta y_t + \gamma z_t < 0 \end{cases}$$

and $\hat{y}_{t+1} = [\alpha + \beta y_t + \gamma z_t]_+$.

6. Empirical Application

It is important to understand the process of the formation of an international trade network of various goods. A good forecast about the future amount of trade is often crucial for numerous policy decisions such as, for example, how much ships/planes/etc. to allocate to the transportation of a given good. The distinguishing feature of trade data is that it is associated with a lot of zeros, as not all countries trade with each other at each moment of time (see e.g. Table 1 in Dueñas and Fagiolo (2013)). This motivates us to use trade data to analyze the performance of the techniques developed in the paper.

We apply our model to monthly exports of pharmaceutical products. The data is obtained from the Eurostat COMEXT database (European Commission (accessed May 1, 2018)). Pharmaceutical industry represents one of the largest industrial sector in the EU and provides a sizable, positive contribution to the EU trade balance (e.g. see Section II in Gambardella et al. (2000) for the discussion based on the Eurostat data). Pharmaceutical industry has a two-digit code 30 as labeled by the Harmonized Commodity Description and Coding Systems (HS). The time period of observation is from January of 1999 until February of 2018, so that $T = 230$. We choose 12 European Union (EU) countries. Those are the countries which

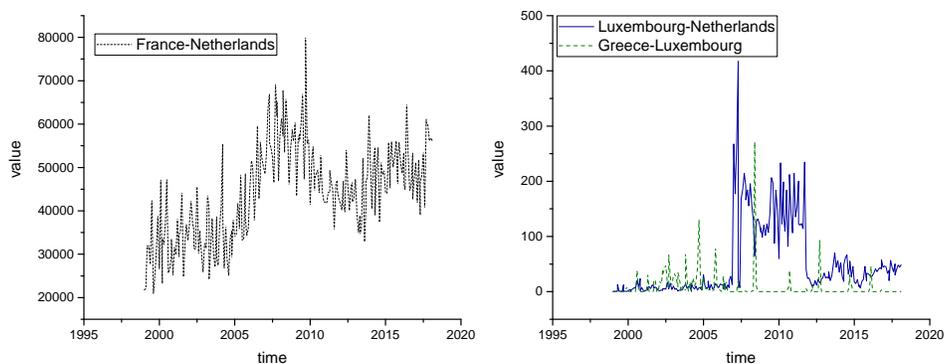


FIGURE 3. Exports of pharmaceutical products between EU countries. Value is in 1000 Euro.

joined the EU first. Thus, each pair of countries represents an edge in the network, and we apply the techniques from previous sections to each such edge.

Figure 3 shows exports of pharmaceutical products between France and Netherlands, Luxembourg and Netherlands, and Greece and Luxembourg. We can see that although the first graph does not have zeros, the second has periods of no trade for both pairs of countries. Zero trade is consistent with our model. The longer length of zero intervals (green dashed curve vs. blue straight one in Figure 3) could be explained by more negative α_{ij} .

To measure predictive power we use a rolling window of size 200. Our consistency results rely on $T \rightarrow \infty$, thus, we want to have as large as possible window size. Additionally, we need more than one window to see how good are our predictions. This leads us to the choice of window size $T' = 200$, which is large enough and still can be rolled enough times⁷. The increment between successive rolling windows is 1 period. For each window we make a 1-period-ahead forecast. That is, based on observations from t to $t + 199$ we predict what happens at $t + 200$ for $t = 1, \dots, 30$. This gives us 30 overlapping windows and, thus, 30 forecasts. We report the average predictive absolute error over them. For testing in the next subsection we need more windows and with that aim in mind we also report results for a rolling window of size 115, which gives 115 overlappind windows (from t to $t + 114$ for $t = 1, \dots, 115$). The reason is that successive windows are very correlated, and to do testing we want to diminish the effect from correlation. Thus, we need more windows over which the test is going to average. In other words, there is a trade-off between maximizing ⁷ $T' = 195$ or $T' = 205$ would be as good and the precise value of T is ad hoc in this sense.

Model and method	R_{abs} in billion euros	
	window size $T' = 200$	window size $T' = 115$
LAD w. p.e.	1.8934	1.8315
MLE w. p.e.	1.9215	1.8274
OLS w. p.e.	1.9217	1.8279
LAD w/o p.e.	1.9865	1.8977
MLE w/o p.e.	2.0036	1.8847
OLS w/o p.e.	2.0037	1.8848
“today”	2.0316	1.9096

TABLE 1. Prediction errors with and without peer effects under different estimation techniques. Smaller numbers mean better quality of prediction. $R_{abs} = \frac{1}{T-T'} \sum_{i,j} \sum_{t=1}^{T-T'} |y_{ij,t+T'} - \hat{y}_{ij,t+T'}|$, where $y_{ij,t+T'}$ is the true value and $\hat{y}_{ij,t+T'}$ is the out-of-sample prediction.

the window size to improve the quality of estimates and maximizing the number of windows (hence, decreasing the size of each) to improve the statistical significance of testing. Without any additional information on this trade-off, we choose the values of two parameters to be equal, leading to window size $T' = 115$. It is impossible to efficiently present the results for all 132 edges individually, so instead we sum over all edges and report the total result.

We consider three alternative benchmark predictions to evaluate the performance of our model and estimation and forecasting techniques. The first benchmark is the naïve forecast “today equals tomorrow”, i.e., $\hat{y}_{t+1} = y_t$. The second alternative is to ignore positivity and treat the model as linear, i.e., $y_t = \alpha + \beta y_{t-1} + \gamma z_{t-1} + u_t$ (OLS in Table 1). Finally, the third competing option is to stick to our model, but to use Gaussian likelihood instead of the LAD (MLE in Table 1). As is shown in the Supplementary Material, in our setting with positive part MLE is consistent when errors are normally distributed (Theorem B.1), yet when the error distribution is far from normal, MLE can perform poorly.

We compare the above three approaches with the LAD estimator with and without peer effects. Results are shown in Table 1, where “w. p.e” stays for the model with peer effects (regressor z_{t-1} is included) and “w/o p.e” stays for the model without peer effects (regressor z_{t-1} is excluded). Let us concentrate on the first column corresponding to the larger rolling window $T' = 200$. We can see that even if we do not incorporate peer effects (last four rows)

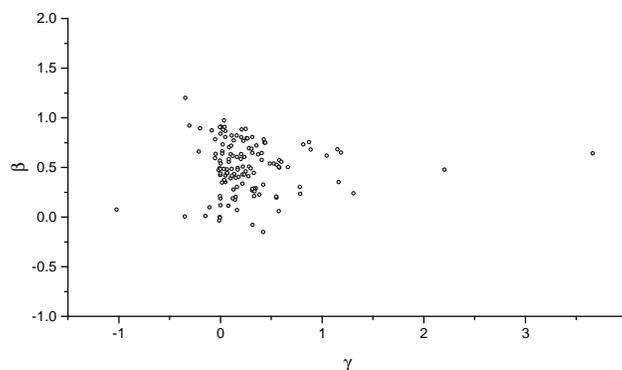


FIGURE 4. LAD estimates for 132 country pairs for the model

$$y_t = [\alpha + \beta y_{t-1} + \gamma z_{t-1} + u_t]_+.$$

model-based predictors outperform all alternatives. The fact that LAD performs significantly better than the MLE suggests that the error distribution may be far from normal.

When we incorporate peer effects, we use the triangular peer effect function: $z_{ijt-1} = p_{ij}(\{y_{k\ell t-1}\}_{k,\ell}) = \sum_k \frac{\sqrt{y_{ikt-1}y_{kjt-1}}}{n-2}$. The first column of Table 1 shows that adding peer effects reduces the prediction error both under LAD and MLE estimation approaches. Same happens if we ignore positivity, assume linear model, and estimate it by OLS (first three rows vs. next three rows). This suggests the presence of peer effects in the data. The scatter plot of the LAD estimates of $(\beta_{ij}, \gamma_{ij})$ is shown in Figure 4. The mean value of β is 0.5037 and the mean value of γ is 0.2652.

In general, including irrelevant regressors in the model leads to more noise when one does prediction. Yet, one can expect that the use of additional relevant regressors may improve the prediction results reported in Table 1. As both columns of Table 2 show, using four lags and one peer effect regressor or using five lags gives the best results ($z_{ijt-1} = \sum_k \frac{\sqrt{y_{ikt-4}y_{kjt-4}}}{n-2}$). Also using peer effects evaluated at $t-4$ leads to better performance than using peer effects evaluated at $t-1$. This suggests that peer effects can help to predict the future, though the optimal functional form of the peer effect function is unclear.

Finally, let us discuss the last columns in Tables 1 and 2 corresponding to the smaller window size $T' = 115$. We see that all methods behave much worse than for $T' = 200$: the improvement of R_{abs} upon the prediction “today equals tomorrow” shrinks significantly. This suggests that the available T is at the border of applicability of our methods.

Estimated model	R_{abs} in billion euros	
	window size $T' = 200$	window size $T' = 115$
5 lags	1.6731	1.6650
4 lags+p.e. at $t - 4$	1.6733	1.6437
4 lags+p.e. at $t - 1$	1.6797	1.6511
1 lag+p.e. at $t - 4$	1.8569	1.8172
1 lag+p.e. at $t - 1$	1.8934	1.8315
1 lag	1.9865	1.8977
“today”	2.0316	1.9096

TABLE 2. Prediction errors with different lags and peer effects. Smaller numbers mean better quality of prediction. $R_{abs} = \frac{1}{T-T'} \sum_{i,j} \sum_{t=1}^{T-T'} |y_{ij,t+T'} - \hat{y}_{ij,t+T'}|$, where $y_{ij,t+T'}$ is the true value and $\hat{y}_{ij,t+T'}$ is the out-of-sample prediction.

6.1. Tests. When we were analyzing the performance of our techniques and comparing predictive errors (Tables 1 and 2), we averaged over all edges of the networks. In the current subsection, we look at how good the prediction is for each given edge. For each edge we compute the Diebold-Mariano test (Diebold and Mariano (1995)) to compare the predictive accuracy under different model specifications. Here we use the rolling window of size 115, as the test requires a large number of predictions. This gives us 115 overlapping windows (from t to $t + 114$ for $t = 1, \dots, 115$). We consider 1-period-ahead forecasts under various models. We compare the model without peer effects against the benchmark prediction “today=tomorrow” and against the model with peer effects. There are 132 edges in total, and generally for between one third and one half of them we reject the null of identical predictive accuracy. This is shown in Table 3. Similarly, Table 4 shows the results of the one-sided Diebold-Mariano tests for different model specifications. Almost all of the rejections in the former table correspond to rejections in the latter table. That is, when we compare two possible model specifications, the one with the positive part and peer effects dominates. This strengthens the importance of incorporating positivity and peer effects. On the other hand, testing the model with 5 lags versus the model with peer effects seems to show much less difference in predictability.

Additionally, for the model with 1 lag and peer effects evaluated at time $t - 1$ we calculate t -statistics to analyze significance of the peer effects for each edge. That is, for each edge

	Number of rejections	
	5%	10%
“today” vs. 1 lag	38	47
“today” vs. 1 lag+p.e. at $t - 1$	45	59
“today” vs. 1 lag+p.e. at $t - 4$	53	60
1 lag vs. 1 lag+p.e. at $t - 1$	39	51
1 lag vs. 1 lag+p.e. at $t - 4$	49	60
5 lags vs. 4 lags+p.e. at $t - 1$	16	22
5 lags vs. 4 lags+p.e. at $t - 4$	15	27

TABLE 3. Two-sided Diebold-Mariano test of equal predictability.

	Number of rejections	
	5%	10%
“today” worse than 1 lag	41	51
“today” worse than 1 lag+p.e. at $t - 1$	52	62
“today” worse than 1 lag+p.e. at $t - 4$	57	65
1 lag worse than 1 lag+p.e. at $t - 1$	38	46
1 lag worse than 1 lag+p.e. at $t - 4$	50	57
5 lags worse than 4 lags+p.e. at $t - 1$	11	21
5 lags worse than 4 lags+p.e. at $t - 4$	21	23

TABLE 4. One-sided Diebold-Mariano test.

($i \rightarrow j$) we calculate $\frac{\sqrt{T}\hat{\gamma}_{ij}}{\hat{std}(\gamma_{ij})}$. We use asymptotic variance from Theorem 6 to estimate $\hat{std}(\gamma_{ij})$. The matrix of second moments M_0 is estimated by its sample analogue, while density at zero, $f_{u_{ij}}(0)$, is estimated by uniform kernel. The bandwidth is chosen so that approximately half of the observations are used. On 5% level peer effects are significant for approximately two thirds of edges. If we instead consider peer effects evaluated at time $t - 4$, the number of significant coefficients goes slightly up (92 out of 132 versus 83 out of 132). This reinforces results of the Diebold-Mariano tests reported in Table 4.

7. Extensions

In Section 2 we specified the model in terms of evolution of edges of a network. The specification allows for both directed and undirected networks. Besides, our framework is also suitable for non-negative panels where n is fixed and T is large and we expect some interactions across units. E.g., alcohol consumption by classmates or behavior of various financial

markets. In this case we can view cross-section units (classmates or financial markets) as vertices in a network and analyze their evolution. If $y_{iit} \equiv y_{it}$ represents the evolution of a characteristic of an agent i , such as amount of alcohol consumed by the agent i at day t , we can use our techniques to estimate and predict future y_{it} s. That is, we can analyze how the alcohol consumption of one's peers affects one's decision to drink or how a crush in one markets affects the behavior of other markets. Mathematically, this corresponds to the non-negative multivariate time series $\{y_{it}\}_{i,t}$:

$$(11) \quad y_{it} = \left[\alpha_i + \beta_i y_{it-1} + \gamma_i p_i \left(\left\{ y_{js} \right\}_{\substack{j=1, \dots, n \\ s=t-H, \dots, t-1}} \right) + u_{it} \right]_+,$$

where, as before, $p_i(\cdot)$ is a peer-effect function. When one analyzes financial markets, we may expect that a crush in one of them pushes the others also to go down. Thus, a peer effect function of the form $\min_j y_{jt}$ may be useful.

Furthermore, our setting allows for joint analysis of the evolution of edges and vertices. If we consider both Eq. (3) and Eq. (11) jointly, then we capture both y_{ijt} , $i \neq j$ and $y_{iit} \equiv y_{it}$. Moreover, peer effect functions p_{ij} and p_i can depend both on y_{kls} and y_{kks} with suitable modification of the Assumption 3. In this extended setting our theorems continue to hold.

A special case of Eq. (11) with linear peer-effect function evaluated at time $t - 1$, $\gamma_i p_i(y_{1t-1}, \dots, y_{nt-1}) = \sum_{j \neq i} \lambda_{ij} y_{ijt-1}$ corresponds to a non-negative VAR(1),

$$\mathbf{y}_t = [\mathbf{A} + \mathbf{B}\mathbf{y}_{t-1} + \mathbf{u}_t]_+,$$

where $\mathbf{y}_t = (y_{1t}, \dots, y_{nt})'$, $\mathbf{u}_t = (u_{1t}, \dots, u_{nt})'$, $\mathbf{A} = (\alpha_1, \dots, \alpha_n)'$ and \mathbf{B} is $n \times n$ matrix with β_i on the main diagonal and λ_{ij} off diagonal. Such setting can be useful in modelling nominal interest rates or realized volatilities.

8. Conclusion

This paper presents a novel approach to modeling and estimating networks. Estimation does not require knowledge of the error distribution, thereby making the whole process more attractive to use. Instead of using the variation across individuals, one can use variation across time to identify parameters of the network. In this approach, we treat networks as

multivariate time series. The main advantage is that we allow the realization of each edge today to depend on the whole structure of the graph in the previous time periods, and not only on the properties of two nodes, which are connected by that edge. Moreover, the Markov form of the equations makes them convenient for doing predictions. As the empirical example suggests, our model does, indeed, help to predict the future. Overall, the results confirm that incorporating non-negativity of the dependent variables into the model matters and incorporating peer effects leads to the improved predictive power.

In the future it would be interesting to apply the model to different data sets. Phone call data for a small group of individuals and technology adoption by countries from one another seem like natural candidates.

From a theoretical point of view, it would be interesting to investigate in more depth the discontinuity in the asymptotics in the model without peer effects. The behaviour of the process $y_t = [y_{t-1} + u_t]_+$ differs dramatically from what one gets by shifting α from zero or β from one slightly. Thus, finding a way to unify the cases in the neighbourhood of the point $(\alpha = 0, \beta = 1)$ in the spirit of Phillips (1987) may be helpful from a practical point of view. Yet, it is a challenging problem. The proper scaling limit of such a process is complicated, because it involves the computation of the time the process spends at zero in the limit.

9. Supplementary Material

See “Evolution of Networks: Supplementary Material” for all proofs and additional results on OLS and MLE estimation.

References

- Amemiya, T.**, “Tobit models: A survey,” *Journal of Econometrics*, 1984, *24* (1–2), 3–61.
- Anderson, T.W.**, “On Asymptotic Distributions of Estimates of Parameters of Stochastic Difference Equations,” *The Annals of Mathematical Statistics*, 1959, *30* (3), 676–678.
- Andrews, D.W.K.**, “Non-Strong Mixing Autoregressive Processes,” *Journal of Applied Probability*, 1984, *21* (4), 930–934.

- Blume, L.E., W.A. Brock, S.N. Durlauf, and Y.M. Ioannides**, *Identification of social interactions*, Vol. 1, North-Holland, 2011. In Handbook of social economics.
- Bramoullé, Y., A. Galeotti, and B. Rogers, eds**, *The Oxford handbook of the economics of networks*, Oxford University Press, 2016.
- Cavaliere, G.**, “The Asymptotic Distribution of the Dickey–Fuller Statistic under Nonnegativity Constraint,” *Econometric Theory*, 2004, *20* (4), 808–810.
- Chamberlain, G.**, “Asymptotic efficiency in semi-parametric models with censoring,” *Journal of Econometrics*, 1986, *32* (2), 189–218.
- de Jong, R. and A.M. Herrera**, “Dynamic censored regression and the Open Market Desk reaction function,” *Journal of Business & Economic Statistics*, 2011, *29* (2), 228–237.
- Diebold, F.X. and K. Yilmaz**, “On the network topology of variance decompositions: Measuring the connectedness of financial firms,” *Journal of Econometrics*, 2014, *182* (1), 119–134.
- and ———, *Financial and macroeconomic connectedness: A network approach to measurement and monitoring*, Oxford University Press, 2015.
- and **R.S. Mariano**, “Comparing predictive accuracy,” *Journal of Business & economic statistics*, 1995, *13* (3), 253–63.
- Dueñas, M. and G. Fagiolo**, “Modeling the international-trade network: a gravity approach,” *Journal of Economic Interaction and Coordination*, 2013, *8* (1), 155–178.
- European Commission**, “Eurostat,” <https://ec.europa.eu/eurostat/> accessed May 1, 2018.
- Gambardella, A., L. Orsenigo, and F. Pammolli**, “Global competitiveness in pharmaceuticals: a European perspective,” *MPRA Paper No. 15965*, 2000.
- Gourieroux, C. and A. Monfort**, “Sufficient linear structures: Econometric applications,” *Econometrica*, 1980, *48* (5), 1083–1097.
- Graham, B.S.**, “Homophily and transitivity in dynamic network formation,” *Working paper*, 2016.
- , “An econometric model of network formation with degree heterogeneity,” *Econometrica*, 2017, *85* (4), 1033–1063.

- , “Dyadic regression,” *Working paper*, 2019.
- Hahn, J. and G. Kuersteiner**, “Stationarity and mixing properties of the dynamic Tobit model,” *Economics Letters*, 2010, *107* (2), 105–111.
- Holme, P. and J. Saramäki**, “Temporal networks,” *Physics reports*, 2012, *519* (3), 97–125.
- Jackson, M.O.**, *Social and economic networks*, Princeton University Press, 2010.
- Khan, S. and J.L. Powell**, “Two-step estimation of semiparametric censored regression models,” *Journal of Econometrics*, 2001, *103* (1–2), 73–110.
- Leung, M.P. and H.R. Moon**, “Normal Approximation in Large Network Models,” *arXiv preprint arXiv:1904.11060*, 2019.
- Michel, J. and R.M. de Jong**, “Mixing properties of the dynamic Tobit model with mixing errors,” *Economics Letters*, 2018, *162*, 112–115.
- Palla, G., A.L. Barabási, and T. Vicsek**, “Quantifying social group evolution,” *Nature*, 2007, *446* (7136), 664–667.
- Phillips, P.C.B.**, “Towards a unified asymptotic theory for autoregression,” *Biometrika*, 1987, *74* (3), 535–547.
- Pin, P. and B.W. Rogers**, “Cooperation, punishment and immigration,” *Journal of Economic Theory*, 2015, *160*, 72–101.
- Powell, J.L.**, “Least absolute deviations estimation for the censored regression model,” *Journal of Econometrics*, 1984, *25* (3), 303–325.
- Robinson, D.T. and T.E. Stuart**, “Network effects in the governance of strategic alliances,” *The Journal of Law, Economics, & Organization*, 2006, *23* (1), 242–273.
- Wang, X. and J. Yu**, “Limit theory for an explosive autoregressive process,” *Economics Letters*, 2015, *126*, 176–180.
- Wei, S.X.**, “A Bayesian approach to dynamic Tobit models,” *Econometric Reviews*, 1999, *18* (4), 417–439.
- White, J.S.**, “The Limiting Distribution of the Serial Correlation Coefficient in the Explosive Case,” *The Annals of Mathematical Statistics*, 1958, *29* (4), 1188–1197.
- Withers, C.S.**, “Central limit theorems for dependent variables. I,” *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 1981, *57* (4), 509–534.

(Anna Bykhovskaya) UNIVERSITY OF WISCONSIN-MADISON

E-mail address: `anna.bykhovskaya@wisc.edu`