



Investigating competition in financial markets: a sparse autologistic model for dynamic network data

Brenda Betancourt, Abel Rodríguez & Naomi Boyd

To cite this article: Brenda Betancourt, Abel Rodríguez & Naomi Boyd (2018) Investigating competition in financial markets: a sparse autologistic model for dynamic network data, *Journal of Applied Statistics*, 45:7, 1157-1172, DOI: [10.1080/02664763.2017.1357684](https://doi.org/10.1080/02664763.2017.1357684)

To link to this article: <https://doi.org/10.1080/02664763.2017.1357684>



Published online: 30 Jul 2017.



Submit your article to this journal



Article views: 71



CrossMark

View Crossmark data



Investigating competition in financial markets: a sparse autologistic model for dynamic network data

Brenda Betancourt^{a,b}, Abel Rodríguez^a and Naomi Boyd^c

^aDepartment of Applied Mathematics and Statistics, University of California, Santa Cruz CA, USA;

^bDepartment of Statistical Science, Duke University, Durham, NC, USA; ^cDepartment of Finance, West Virginia University, Morgantown, WV, USA

ABSTRACT

We develop a sparse autologistic model for investigating the impact of diversification and disintermediation strategies in the evolution of financial trading networks. In order to induce sparsity in the model estimates and address substantive questions about the underlying processes the model includes an L^1 regularization penalty. This makes implementation feasible for complex dynamic networks in which the number of parameters is considerably greater than the number of observations over time. We use the model to characterize trader behavior in the NYMEX natural gas futures market, where we find that disintermediation and not diversification or momentum tend to drive market microstructure.

ARTICLE HISTORY

Received 27 June 2016

Accepted 13 July 2017

KEYWORDS

Financial trading networks; autologistic models; Lasso penalty; dynamic networks; link prediction

1. Introduction

Historically, the empirical study of financial markets has emphasized the behavior of aggregate measures such as price discovery or transaction volumes. However, the challenges that come from the rise of electronic and automated trading have highlighted the need to study patterns of individual market transactions (the so-called market microstructure) in order to understand the mechanisms that underpin price formation.

Financial Trading Networks (FTNs), which are directed graphs in which nodes correspond to traders operating in a financial market and edges/weights represent pairwise buy-sell transactions among them that occur within a period of time, are becoming a popular tool for studying the complexity associated with modern financial markets. Indeed, FTNs contain key information about patterns of order execution in order-driven markets, which can in turn provide important insights into the functioning of the market. Empirical work on trading networks so far has focused on studying the evolution of summary statistics such as degree distributions, average betweenness and clustering coefficients and their relationship with market variables such as the volatility of returns (e.g. see [1]). Model-based approaches are quite rare, one exception is Betancourt *et al.* [4], where a hidden Markov model for dynamic network data is introduced to identify change points in the underlying market microstructure.

In this paper we are interested in using FTNs to investigate whether traders in financial markets engage in strategic behaviors such as diversification and disintermediation as part of their long term trading strategies. To accomplish this goal we extend the notion of an autologistic model to directed binary network data. In the much simpler case of a binary time series y_1, y_2, \dots, y_T with $y_t \in \{0, 1\}$, a first-order autologistic model with parameters α and ξ assumes that

$$\text{logit } \Pr(y_t = 1 \mid y_{t-1}, \dots, y_1) = \eta_t = \alpha + \xi y_{t-1}, \quad (1)$$

where $\text{logit } x = \log\{x/(1-x)\}$ and the unknown parameters α and ξ control the structure of the temporal dependence (in particular, note that $\xi = 0$ implies that the observations are independent and identically distributed). This implies that

$$p(y_2, \dots, y_T \mid y_1, \alpha, \xi) = \prod_{t=2}^T \frac{\exp\{y_t(\alpha + \xi y_{t-1})\}}{1 + \exp\{\alpha + \xi y_{t-1}\}}.$$

Autologistic models for spatio-temporal binary data have been discussed in [33,34]. However, these models for spatio-temporal data cannot be directly applied to the network time series data discussed in this paper because they are not designed to account for common features of directed network data such as reciprocity (e.g. the tendency of nodes in the network to consistently respond to a positive action with another positive action) and transitivity (e.g. the tendency of nodes to interact if they share links with a common third party). In contrast, the class of models we introduce in this paper are specifically designed to account for these features, and its parameters have a direct interpretation in terms of network properties.

The autologistic models developed in this paper are special cases of the so-called p_1 models of Holland and Leinhardt [14], which have been extended to dynamic settings in [3,10,18], among others. However, our approach differs from these by directly incorporating the observed network links at previous time points rather than relying on summary network statistics. Our approach is also loosely related to the temporal version of the Exponential Random Graph model (ERGM) introduced in [12] and further developed in [7,26]. However, our approach allows for networks effects to be different for each pair of nodes, leading to additional expressive power and richer interpretation. Other relevant approaches for modeling network data include the dynamic version of the latent space model of Hoff *et al.* [13] developed by Sarkar and Moore [23] and Sewell and Chen [24], the work of Xing *et al.* [31] presenting the temporal extension of the stochastic blockmodel first introduced in [2], the work of Huang and Lin [15], who present an autoregressive integrated moving average model and combine it with link occurrence scores based on similarity indices of network topology measures (e.g. Adamic-Adar coefficient, Katz index), and Bliss *et al.* [6], who propose a method based on similarity indices and node attributes (e.g. common neighbors and preferential attachment) together with a covariance matrix adaptation evolution strategy for link prediction in networks with a large number of nodes.

One challenge for statistical inference for the class of models we discuss in this paper is that the number of parameters is large and grows linearly with the number of nodes in the network. Indeed, unless the system is observed for a very long time, the number of parameters in our autologistic model will typically be much larger than the number of

available observations. To deal with this challenge we employ L^1 regularization which is particularly useful to reduce the number of parameters in the $p > N$ case. The use of an L^1 penalty leads to sparse solutions in which a large number of model coefficients are set to zero, allowing us to address substantive questions about the type of processes driving the evolution of the network. In addition, while performing short-term predictions is a secondary goal for our model, prediction accuracy can sometimes be improved by shrinking the model coefficients [27]. Computation of the procedure is carried out using a coordinate descent method on a surrogate quadratic approximation for an L^1 regularized multinomial likelihood (e.g. see [9]).

The methodology described in this paper is illustrated using a dataset on transactions in the NYMEX natural gas futures market that took place between January 2005 and December 2008. This dataset was previously analyzed in [4] using a hidden Markov model to identify points of structural change in the market. In contrast, the analysis in this paper suggests that disintermediation effects tend to be the most important drivers of network evolution, a pattern that was previously unknown and is consistent with competitive markets.

The remainder of the paper is organized as follows: Section 3 describes our model and discusses some of its properties. Section 4 describes our computational algorithm and some of the properties of the estimators. Section 5 discusses two illustrations, one based on simulated data and a second one that focuses on real trading networks from the natural gas futures market on the New York Mercantile Exchange (NYMEX). Finally, Section 6 presents a short discussion.

2. DATA

The dataset we analyze in this paper consists of individual transaction records from the New York Mercantile Exchange (NYMEX), a commodity futures exchange owned by the Chicago Mercantile Exchange. Commodities traded on NYMEX include coal, electricity, palladium, uranium, and natural gas, among others. Our analysis focuses on proprietary trades (i.e. transactions carried out by traders for their own accounts rather than on their client's behalf) in the natural gas futures market covering the period from January 2005 to December 2008. During this period, the NYMEX natural gas market operated as an open outcry market until 5 September 2006, and as a hybrid market that included electronic trading (conducted via the CME Globex platform) after that date. Over 900 unique traders participated in the market during the period under study; however, the vast majority of the traders participated in the market only sporadically. Furthermore, we have no access to information about whether a specific trader entered or left the market at a given point in time. Hence, we focus our analysis on 71 large traders identified as being present in the market (although not necessarily active) during the whole period. From the original transaction data we construct a sequence of weekly binary FTNs by setting the entries of the adjacency matrices to $y_{i,j,t} = 1$ if there was at least one transaction in which trader i sold a contract to trader j during week t .

3. Modeling approach

Consider a sequence of T binary directed networks, each one observed over a common set of n nodes. The adjacency matrix of the network at time t is therefore an $n \times n$ binary

matrix $\mathbf{Y}_t = [y_{i,j,t}]$, where $y_{i,j,t} = 1$ if there is a link directed from node i to node j at time t , and $y_{i,j,t} = 0$ otherwise. We adopt the convention $y_{i,i,t} \equiv 0$ so that there are no loops within the network. In the illustration we discuss in Section 5, the nodes in the network correspond to traders in NYMEX natural gas futures market, so that $y_{i,j,t} = 1$ if trader i sold a contract to trader j at least once during week t .

We consider an extension of (1) in which the pairs $\{(y_{i,j,t}, y_{j,i,t}) : i < j\}$ are assumed conditionally independent given the history of the network, and each pair $(y_{i,j,t}, y_{j,i,t})$ is modeled using a logistic model of the form

$$p(y_{i,j,t}, y_{j,i,t} \mid \mathbf{Y}_{t-1}) = \exp\{\eta_{i,j,t,1}y_{i,j,t} + \eta_{i,j,t,2}y_{j,i,t} + \eta_{i,j,t,3}y_{i,j,t}y_{j,i,t} - C(\eta_{i,j,t,1}, \eta_{i,j,t,2}, \eta_{i,j,t,3})\}, \quad (2)$$

where the normalizing factor is given by

$$C(\eta_{i,j,t,1}, \eta_{i,j,t,2}, \eta_{i,j,t,3}) = \log[1 + \exp\{\eta_{i,j,t,1}\} + \exp\{\eta_{i,j,t,2}\} + \exp\{\eta_{i,j,t,1} + \eta_{i,j,t,2} + \eta_{i,j,t,3}\}],$$

and the natural parameters $\eta_{i,j,t,1} = f_{i,j,1}(\mathbf{Y}_{t-1})$, $\eta_{i,j,t,2} = f_{i,j,2}(\mathbf{Y}_{t-1})$ and $\eta_{i,j,t,3} = f_{i,j,3}(\mathbf{Y}_{t-1})$ depend on time only through \mathbf{Y}_{t-1} . Note that $\eta_{i,j,t,3}$ controls the level of dependence between $y_{i,j,t}$ and $y_{j,i,t}$. For example, $\eta_{i,j,t,3} = 0$ implies that $y_{i,j,t}$ and $y_{j,i,t}$ are conditionally independent with $\Pr(y_{i,j,t} = 1 \mid \mathbf{Y}_{t-1}) = \exp\{\eta_{i,j,t,1}\}/(1 + \exp\{\eta_{i,j,t,1}\})$ and $\Pr(y_{j,i,t} = 1 \mid \mathbf{Y}_{t-1}) = \exp\{\eta_{i,j,t,2}\}/(1 + \exp\{\eta_{i,j,t,2}\})$. On the other hand, $\eta_{i,j,t,3} > 0$ favors outcomes in which $y_{i,j,t} = y_{j,i,t}$ (a phenomenon often called positive reciprocity), while $\eta_{i,j,t,3} < 0$ favors situations in which $y_{i,j,t} \neq y_{j,i,t}$ (often called negative reciprocity). Hence, by allowing the values of $y_{i,j,t}$ and $y_{j,i,t}$ to be potentially correlated the model can accommodate (intra-temporal) reciprocity in the network.

A full specification of the model requires that we specify the form of the functions $f_{i,j,1}$, $f_{i,j,2}$ and $f_{i,j,3}$. A tempting option is to make these predictors dependent of all entries of \mathbf{Y}_{t-1} , including all high-order interactions. However, such an approach leads to models with an extremely high number of parameters that is computationally unmanageable even for networks with a relatively small number of nodes. On the other hand, while focusing only on first-order effects associated with the entries of \mathbf{Y}_{t-1} can substantially reduce the number of parameters, the resulting model ignores interactions that could be expected to be important. We take a middle ground approach and include in the specification of the functions $f_{i,j,1}$, $f_{i,j,2}$ and $f_{i,j,3}$ a subset of the first- and second-order effects that are associated with the interactions of nodes i and j among themselves and with other nodes during the previous period. In particular, we set

$$\begin{aligned} f_{i,j,l}(\mathbf{Y}_{t-1}) = & \alpha_{i,j,l} + \beta_{i,j,l}y_{i,j,t-1} + \gamma_{i,j,l}y_{j,i,t-1} \\ & + \sum_{k \neq i,j} \delta_{i,j,k,l}y_{i,k,t-1} + \sum_{k \neq i,j} \phi_{i,j,k,l}y_{k,j,t-1} \\ & + \sum_{k \neq i,j} \psi_{i,j,k,l}y_{j,k,t-1} + \sum_{k \neq i,j} \omega_{i,j,k,l}y_{k,i,t-1} \end{aligned}$$

$$\begin{aligned}
& + \rho_{i,j,l} y_{i,j,t-1} y_{j,i,t-1} + \sum_{k \neq i,j} \xi_{i,j,k,l} y_{i,k,t-1} y_{k,j,t-1} \\
& + \sum_{k \neq i,j} \zeta_{i,j,k,l} y_{j,k,t-1} y_{k,i,t-1}
\end{aligned} \tag{3}$$

for $l = 1, 2, 3$. To better motivate this specification, consider for example the structure of $f_{i,j,1}(\mathbf{Y}_{t-1})$ in Equation (3). As we showed before, we can roughly interpret $f_{i,j,1}(\mathbf{Y}_{t-1})$ as controlling the probability of a directed link between i and j . Hence, $\alpha_{i,j,1}$ can be interpreted as the baseline probability of a link between nodes i and j , the coefficients $\beta_{i,j,1}$ and $\gamma_{i,j,1}$ can be interpreted as the persistence (momentum) in the relationship (e.g. if $\beta_{i,j,1} > 0$ then once trader i starts selling to trader j , they tend to keep selling in future periods), the coefficients $\{\delta_{i,j,k,1} : k \neq i,j\}$, $\{\phi_{i,j,k,1} : k \neq i,j\}$, $\{\psi_{i,j,k,1} : k \neq i,j\}$ and $\{\omega_{i,j,k,1} : k \neq i,j\}$ capture diversification effects (e.g. if $\delta_{i,j,k,1} > 0$ then it is more likely that i will sell to j if it sold to k in the previous term), $\rho_{i,j,1}$ captures inter-temporal reciprocity (as opposed to the intra-temporal reciprocity captured by $\eta_{i,j,t,3}$), and $\{\xi_{i,j,k,1} : k \neq i,j\}$ and $\{\zeta_{i,j,k,1} : k \neq i,j\}$ capture disintermediation effects (e.g. if $\xi_{i,j,k,1} > 0$ then so that i is more likely to sell to j if in the previous period i sold to k and k sold to j , so that i and j tend to cut k as middleman).

Although our model is not in the class of time-varying Exponential Random Graph models (tERGMs) [7, 12, 26], some classes of tERGMs can be obtained as special cases of our model. Indeed, consider making the model parameters independent of the traders' identities so that $\alpha_{i,j,l} = \alpha_l$, $\beta_{i,j,l} = \beta_l$, $\gamma_{i,j,l} = \gamma_l$, $\delta_{i,j,k,l} = \delta_l$, $\phi_{i,j,k,l} = \phi_l$, etc. for all i,j,k . In that case, the joint distribution $p(\mathbf{Y}_2, \dots, \mathbf{Y}_T | \mathbf{Y}_1)$ is proportional to

$$\begin{aligned}
& \exp \left\{ \sum_{t=2}^T \sum_{l=1}^3 [\alpha_l S_{\alpha,l}(\mathbf{Y}_t) + \beta_l S_{\beta,l}(\mathbf{Y}_t, \mathbf{Y}_{t-1}) \right. \\
& + \gamma_l S_{\gamma,l}(\mathbf{Y}_t, \mathbf{Y}_{t-1}) + \delta_l S_{\delta,l}(\mathbf{Y}_t, \mathbf{Y}_{t-1}) \\
& + \phi_l S_{\phi,l}(\mathbf{Y}_t, \mathbf{Y}_{t-1}) + \psi_l S_{\psi,l}(\mathbf{Y}_t, \mathbf{Y}_{t-1}) \\
& + \omega_l S_{\omega,l}(\mathbf{Y}_t, \mathbf{Y}_{t-1}) + \rho_l S_{\rho,l}(\mathbf{Y}_t, \mathbf{Y}_{t-1}) \\
& \left. + \xi_l S_{\xi,l}(\mathbf{Y}_t, \mathbf{Y}_{t-1}) + \zeta_l S_{\zeta,l}(\mathbf{Y}_t, \mathbf{Y}_{t-1})] \right\}, \tag{4}
\end{aligned}$$

where $S_{\alpha,l}(\mathbf{Y}_t, \mathbf{Y}_{t-1})$, $S_{\beta,l}(\mathbf{Y}_t, \mathbf{Y}_{t-1})$, etc. are appropriately chosen sufficient statistics, for example,

$$\begin{aligned}
S_{\alpha,1}(\mathbf{Y}_t) &= \sum_{i=1}^n \sum_{j=i+1}^n y_{i,j,t}, \\
S_{\alpha,2}(\mathbf{Y}_t) &= \sum_{i=1}^n \sum_{j=1}^{i-1} y_{i,j,t}, \\
S_{\alpha,3}(\mathbf{Y}_t) &= \sum_{i=1}^n \sum_{j=1}^{i-1} y_{i,j,t} y_{j,i,t},
\end{aligned}$$

$$S_{\delta,2}(\mathbf{Y}_t, \mathbf{Y}_{t-1}) = \sum_{i=1}^n \sum_{j=1}^{i-1} y_{i,k,t-1} y_{i,j,t}$$

$$S_{\xi,1}(\mathbf{Y}_t, \mathbf{Y}_{t-1}) = \sum_{i=1}^n \sum_{j=i+1}^n \sum_{k \neq i,j} y_{i,j,t} y_{i,k,t-1} y_{k,j,t-1}.$$

Hence, if we were to strip away the identity of the nodes in the definition of the model coefficients, the model would reduce to a tERGM constructed on the basis of sufficient statistics that correspond to the number of links in the network as well as the number of (some selected types of) two-stars and triangles. By allowing the parameters to differ according to the identity of the nodes, our formulation generalizes the basic tERGM and allows for additional expressive power.

It is worthwhile noting that the collapsed model in Equation (4) does not include triangles in which all observations happen in the same time point, for example, $\sum_{i=1}^n \sum_{j=i+1}^n \sum_{k \neq i,j} y_{i,j,t} y_{i,k,t} y_{k,j,t}$. Hence, our model cannot capture the effects of *intra-temporal* transitivity (i.e. an increase/decrease in the probability of a link between nodes i and j at time t if they both link to a third node k also at time t) on the evolution of the network. This modeling choice is made out of practical necessity; including this type of interactions into the model would complicate computation. Indeed, assuming conditional independence among pairs of dyads on the same point in time is key to obtain a closed-form structure for the normalizing constant of the likelihood, which is in turn key to speed-up computation (see Section 4). However, potential concerns surrounding this choice are mitigated by the fact that the second-order interactions in Equation (3) do allow us to capture *inter-temporal* transitivity (i.e. an increase/decrease in the probability of a link between nodes i and j at time t if they both linked to a third node k at time $t-1$), which is more interesting and realistic in this type of scenarios.

3.1. A penalized regression model

The total number of parameters in our model is $(n(n-1)/2)\{9 + 18(n-2)\}$, which will typically be quite large. In fact, the number of parameters in the model will often be larger than the number of observations available to estimate them. To address this issue we adopt a regularized likelihood approach based on L^1 penalty functions. More specifically, point estimates for the model parameters are obtained by solving:

$$\arg \max_{\boldsymbol{\alpha}, \boldsymbol{\Theta}} \sum_{i < j} \{V_{i,j}(\boldsymbol{\alpha}_{i,j}, \boldsymbol{\Theta}_{i,j}) - \lambda \|\boldsymbol{\Theta}_{i,j}\|_1\}, \quad (5)$$

where

$$V_{i,j}(\boldsymbol{\alpha}_{i,j}, \boldsymbol{\Theta}_{i,j}) = \sum_{t=2}^T \{y_{i,j}[\alpha_{i,j,1} + \mathbf{x}_{i,j,t}^T \boldsymbol{\theta}_{i,j,1}] + y_{j,i}[\alpha_{i,j,2} + \mathbf{x}_{i,j,t}^T \boldsymbol{\theta}_{i,j,2}]$$

$$+ y_{i,j}y_{j,i}[\alpha_{i,j,3} + \mathbf{x}_{i,j,t}^T \boldsymbol{\theta}_{i,j,3}]\}$$

is the (unpenalized) log-likelihood, $\|\cdot\|_1$ denotes the L_1 -norm, $\lambda > 0$ is the penalty parameter that controls the shrinkage level of the coefficients towards zero, $\boldsymbol{\alpha}_{i,j} =$

$(\alpha_{i,j,1}, \alpha_{i,j,2}, \alpha_{i,j,3})'$, the vector of parameters $\theta_{i,j,r}$ is defined as

$$\begin{aligned}\theta_{i,j,r} = & (\beta_{i,j,r}, \gamma_{i,j,r}, \delta_{i,j,1,r}, \dots, \delta_{i,j,n,r}, \phi_{i,j,1,r}, \dots, \phi_{i,j,n,r}, \psi_{i,j,1,r}, \dots, \psi_{i,j,n,r}, \omega_{i,j,1,r}, \\ & \dots, \omega_{i,j,n,r}, \rho_{i,j,r}, \xi_{i,j,1,r}, \dots, \xi_{i,j,n,r}, \zeta_{i,j,1,r}, \dots, \zeta_{i,j,n,r})',\end{aligned}$$

$\Theta_{i,j} = (\theta'_{i,j,1}, \theta'_{i,j,2}, \theta'_{i,j,3})'$, and the vector of covariates is

$$\begin{aligned}\mathbf{x}_{i,j,t} = & (y_{i,j,t-1}, y_{j,i,t-1}, y_{i,1,t-1}, \dots, y_{i,n,t-1}, y_{1,j,t-1}, \\ & \dots, y_{n,j,t-1}, y_{j,1,t-1}, \dots, y_{j,n,t-1}, y_{1,i,t-1}, \\ & \dots, y_{n,i,t-1}, y_{i,j,t-1} y_{j,i,t-1}, y_{i,1,t-1} y_{1,j,t-1}, \\ & \dots, y_{i,n,t-1} y_{n,j,t-1}, y_{j,1,t-1} y_{1,i,t-1}, \dots, y_{j,n,t-1} y_{n,i,t-1})'.\end{aligned}$$

Note that the structure $V_{i,j}$ is equivalent to that of a multinomial likelihood and that the intercept parameters $\{\alpha_{i,j,r}\}$ remain unpenalized. Furthermore, the imposition of a lasso penalty is equivalent to assuming independent double exponential prior distributions with variance $2/\lambda$ on each component of Θ , so that the point estimates obtained from Equation (5) coincide with the maximum a posteriori estimates.

4. Estimation and prediction

One important consequence of the conditional independence assumption is that Equation (5) can be broken down into $n(n - 1)/2$ optimization problems, each one corresponding to fitting a separate L^1 regularized multinomial regression for each pair of nodes in the network. In the sequel we focus on an algorithm to solve each of these independent problems (and drop the subindex (i,j) to simplify notation). This algorithm is then implemented in a parallel environment.

There is an extensive literature on efficient algorithm for estimation in L^1 regularized multinomial regression (e.g. see [9,11,19]). In this paper we resort to a relatively simple computational algorithm similar to iterative reweighted least squares (see [9]). In particular, we solve Equation (5) by iteratively setting

$$(\hat{\alpha}^{(m+1)}, \hat{\Theta}^{(m+1)}) = \underset{\Theta}{\operatorname{argmax}} Q(\alpha, \Theta \mid \hat{\alpha}^{(m)}, \hat{\Theta}^{(m)})$$

until convergence, where $Q(\alpha, \Theta \mid \tilde{\alpha}, \tilde{\Theta})$ is a surrogate function obtained by replacing $V(\alpha, \Theta)$ by its second-order Taylor expansion around the current iterate. Furthermore, rather than attempting to solve the problem using blockwise updates, we proceed with componentwise steps. In particular, the estimate of a component $\theta_{r,k}$ of θ_r is updated as

$$\hat{\theta}_{r,k}^{(m+1)} = \operatorname{soft} \left(\hat{\theta}_{r,k}^{(m)} - \frac{g_{r,k}^{(m)}}{G_{r,k,k}^{(m)}}; \frac{-\lambda}{G_{r,k,k}^{(m)}} \right), \quad (6)$$

where $\operatorname{soft}(w, \lambda) = \operatorname{sign}(w) \max\{0, |w| - \lambda\}$ is the soft-thresholding operator, $g_{r,k}^{(m)} = \partial V / \partial \theta_{r,k} \big|_{\hat{\alpha}^{(m)}, \hat{\Theta}^{(m)}}$ and $G_{r,k,k}^{(m)} = -\partial^2 V / \partial \theta_{r,k}^2 \big|_{\hat{\alpha}^{(m)}, \hat{\Theta}^{(m)}}$ are the gradient and the information

in the direction of $\theta_{r,k}$ evaluated in the current iterate values. On the other hand, since the intercepts are not penalized, their estimates are updated using the recursion

$$\hat{\alpha}_k^{(m+1)} = \hat{\alpha}_k^{(m)} - \frac{g_k^{(m)}}{G_{k,k}^{(m)}}.$$

One consequence of the use of L_1 penalized likelihoods is that point estimates of the coefficients can be exactly zero. Hence, the algorithm automatically performs variable selection, allowing us to assess the presence of diversification and disintermediation effects in the network. This allows us to explicitly test hypotheses about the kind of effects that influence the evolution of the network. However, when the regression matrix is not full rank (for example, when $T < 9 + 18(n - 2)$), interpretation of the individual effects is difficult because of confounding/multicollinearity. To address this issue we focus on identifying regression coefficients for which there is no evidence of significance. These are selected by identifying the effects that lie in the orthogonal complement of the column space of $\mathbf{X}_{i,j}(\mathcal{A}_\lambda)$, the submatrix that contains the columns associated with variables that have been identified as significant using the penalty λ .

4.1. Selection of the penalty parameter

The value of the penalty λ has a direct impact on the quality of the estimates and predictions generated by the model. Our default approach is to select λ from among a pre-specified grid of values by maximizing the Bayesian Information Criteria (BIC)

$$\text{BIC}_\lambda = \sum_{i < j} [2V_{i,j}(\hat{\alpha}_{i,j}, \hat{\Theta}_{i,j}) - \mathcal{K}_{i,j}(\lambda) \log(T - 1)],$$

where $\mathcal{K}_{i,j}(\lambda) = \text{rank}\{\mathbf{X}_{i,j}(\mathcal{A}_\lambda)\}$ is an estimate of the number of degrees of freedom when the penalty parameter λ is used to compute $(\hat{\alpha}_{i,j}, \hat{\Theta}_{i,j})$, and $\mathbf{X}_{i,j}(\mathcal{A}_\lambda)$ is a $(T - 1) \times d$ matrix whose t -th row contains a subset of elements of $\mathbf{x}_{i,j,t}$ and whose columns correspond to the covariates for which $\hat{\theta}_{i,j,r}$ is different from zero for at least one value $r = 1, 2, 3$ [22, 28, 36]. Note that for all values of λ , the degrees of freedom satisfy the condition $0 \leq \mathcal{K}_\lambda \leq \min\{d, T - 1\}$.

4.2. Link prediction

Although the main goal of our analysis is not short-term link prediction, our autologistic model can be used for this purpose. In particular, given a point estimate $(\hat{\alpha}_{i,j}, \hat{\Theta}_{i,j})$ based on an observed sample $\mathbf{Y}_1, \dots, \mathbf{Y}_T$, we can estimate (for $i > j$) the probability of a directed link from node i to node j at time $T+1$ as

$$\begin{aligned} \hat{p}(y_{i,j,T+1} = 1 \mid \mathbf{Y}_T) &= p[(y_{i,j,T+1}, y_{j,i,T+1}) = (1, 0) \mid \hat{\alpha}_{i,j}, \hat{\Theta}_{i,j})] \\ &\quad + p[(y_{i,j,T+1}, y_{j,i,T+1}) = (1, 1) \mid \hat{\alpha}_{i,j}, \hat{\Theta}_{i,j})], \end{aligned}$$

with a similar expression being valid for $\hat{p}(y_{j,i,T+1} = 1 \mid \mathbf{Y}_T)$.

4.3. Theoretical properties

Because our estimation procedure reduces to fitting independent L^1 regularized multinomial logistic regressions for each pair of nodes, the procedure shares all the positive (and negative) properties of this type of approaches (for a review, see [29]). For example, for n fixed and T growing to infinity, the estimators are both consistent and sparse consistent, and have the oracle property (see [8] for the full rank case and [20,32] for the non-full rank case). Under some additional restrictions (such as the Restricted Eigenvalue (RE) condition described in [5]), these results apply for n growing with T (e.g. see [17,30]).

5. Application

5.1. Simulation study

We begin by demonstrating the predictive performance of the model in a simulated dataset consisting of $T = 201$ networks observed over $n = 71$ nodes (the same values as in the NYMEX data). The data was generated according to our model in such a way that for each of the 2,485 pairs of nodes only six non-zero coefficients are present for each class, $l = 1, 2, 3$. Three of the non-zero coefficients for each class correspond to $\alpha_{i,j,l}$, $\beta_{i,j,l}$, and $\gamma_{i,j,l}$. We randomly draw these parameters from common Gaussian distributions across pairs (e.g. $\alpha_{i,j,l} \sim N(\bar{\alpha}_l, \tau_l^2)$). The other relevant coefficients correspond to $\xi_{i,j,k,l}$ for three different values of k . Four groups of pairs of traders of similar sizes were simulated with different selections of the three values of k , and the respective parameter values of $\xi_{i,j,k,l}$ where fixed equal within each of the groups and with opposites signs to the global mean of $\beta_{i,j,l}$ and $\gamma_{i,j,l}$. Under this simulation scheme, the persistence of the relationship between the nodes i and j and a few transitive relationships drive the network structure and dynamics over time. The resulting network is relatively dense with a average number of links of 2971 (out of 4970 possible ties), and it shows low reciprocity and high transitivity.

As a benchmark, we also fit to the data a tERGM [21] that includes all the typical ERGM terms, the square root of in and out-degrees as node covariates, and the lagged network and the delayed reciprocity to model cross-temporal dependencies. This model was fit using the `btergm` function of the R package `xergm`. Our evaluation is based on an out-of-sample crossvalidation exercise where we held out the last 10 weeks in the data set and made one-step-ahead predictions for the structure of the held-out networks. More specifically, for each $t = 191, 192, \dots, 200$ we use the information contained in $\mathbf{Y}_1, \dots, \mathbf{Y}_t$ to estimate the model parameters and obtain predictions for $\hat{\mathbf{Y}}_{t+1}$. The quality of the prediction is evaluated by constructing receiver operating characteristic (ROC) curves and computing the area under this curve. Predictions for the tERGM model are based on 1000 MCMC simulations generated using the default parameter values for the `xergm` package (see `btergm` documentation for more details).

We search for the optimal value of λ over a grid of 29 values between 3.5 and 25, with the optimal value being $\lambda = 12$ (see Figure 1). Figure 2 shows the 10 operating characteristic curves associated with one-step-ahead out of sample predictions from our autologistic model, along with estimates of the area under the receiver operating characteristic curves (AUC) for the proposed model and the tERGM. From these results it can be seen that the predictive accuracy of the temporal ERGM is poor as it only reaches AUC values below

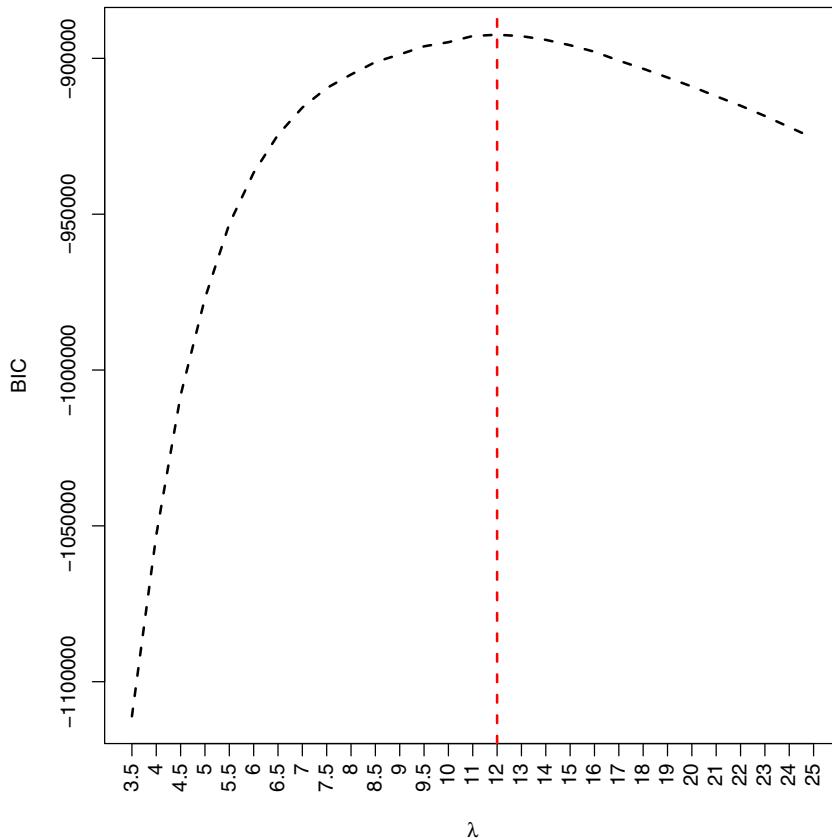


Figure 1. BIC values over a grid of values of λ for autologistic model in simulated dataset.

75% in most cases. In this scenario, the autologistic model shows superior prediction ability outperforming the tERGM by 11% to 14% in the AUC values for all cases.

5.2. Analysis of the NYMEX data

In this section we analyze a sequence of $T = 201$ weekly FTNs constructed from proprietary trades in the natural gas futures market on the New York Mercantile Exchange (NYMEX) between January 2005 and December 2008 (see Section 2). Previous exploration of this data showed that these trading networks are moderately sparse (with an average of 826 links out 4970 possible ones), and consistently show very high reciprocity, moderate transitivity, mixing patterns and community structure [4]. These features suggests that link formation between a pair of nodes is very likely to depend on how other actors relate in the network.

Selection of the optimal penalization parameter in this case was performed by searching over a grid of 24 values between 2.5 and 18 for a resulting optimal value of $\lambda = 10$ (see Figure 3). As before, we carry out an out-of-sample cross-validation exercise in which our model and the tERGM are fitted to the first 191 weeks, and the estimated model parameters are then used to predict each of the last 10 weeks of data. Figure 4 shows the ten operating characteristic curves associated with the out-of-sample predictions from the autologistic

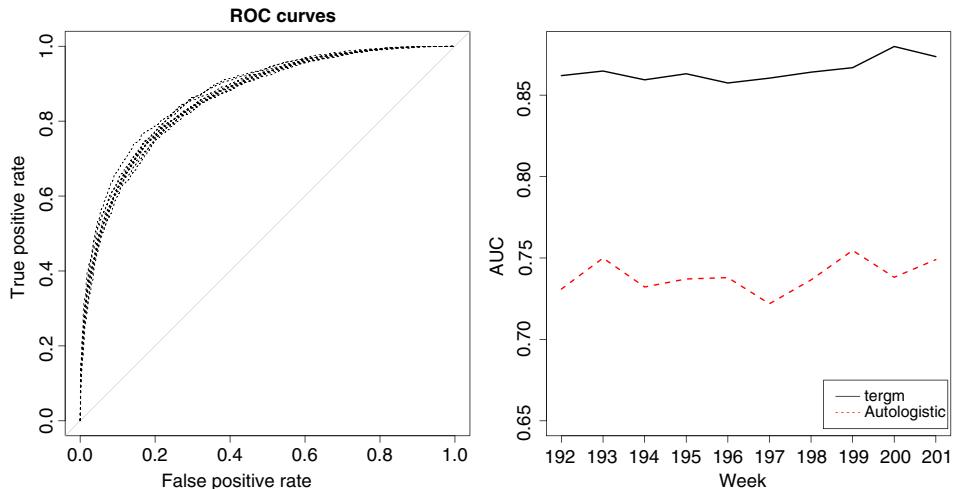


Figure 2. Plots of the 10 operating characteristic curves associated with one-step-ahead out of sample predictions from the autologistic model, and the area under the curves (AUC) for the autologistic model and the temporal ERGM for the simulated dataset.

model, and the estimates of the area under the ROC curves (AUC) for both models. In this particular case, the temporal ERGM outperforms our proposed model by between 3% and 6% in the AUC. However, the prediction accuracy of the autologistic model is reasonably good with an average AUC value of 85% over the 10 weeks. The temporal ERGM slightly outperforming our model for the NYMEX data is probably a result of the presence of moderate *intra-temporal* transitivity that we are unable to capture in our model. In contrast, the simulated network in Section 5.1 is dominated by *inter-temporal* transitivity that the temporal ERGM is unable to capture.

Now, we turn our attention to the interpretation of the regression coefficients. Recall that the trading network has $n = 71$ traders, so that the autologistic model includes 1251 covariates (excluding the unpenalized intercept). Of these 1251 coefficients, 6 capture persistence effects, 3 capture inter-temporal reciprocity, 828 capture substitution effects, and 414 capture disintermediation effects. However, since the number of covariates is much larger than the number of observations, these effects are confounded with each other, complicating the interpretation of the model. To address this issue we focus on identifying effects for which there is no evidence of significance (see Section 4).

First, we note that the individual regression models for each pair tend to be quite sparse. Indeed, only 812 out of 2485 regressions have at least one non-zero regression coefficient aside from the intercept and the number of non-zero coefficients (i.e. the rank of $\mathbf{X}_{i,j}(\mathcal{A}_\lambda)$) tend to be very low across these pairs (see first panel Figure 5). However, the number of significant effects varies dramatically across the different pairs. For example, the interaction between traders 2 and 17 (we identify traders by a number rather than their name because of confidentiality restrictions) seems to be driven by five significant effects: two of them are associated with persistence, one with reciprocity, and the other two with substitution/diversification. This is in contrast with the interaction between traders 64 and 71, which appear to be driven by over 100 potentially significant effects.

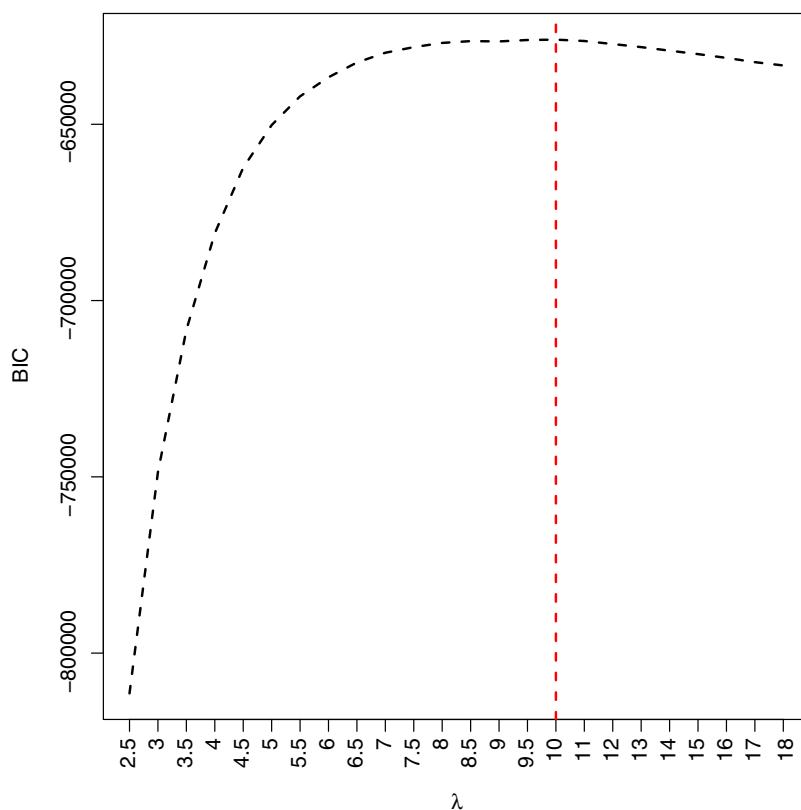


Figure 3. BIC values over a grid of values of λ for autologistic model in NYMEX financial trading network.

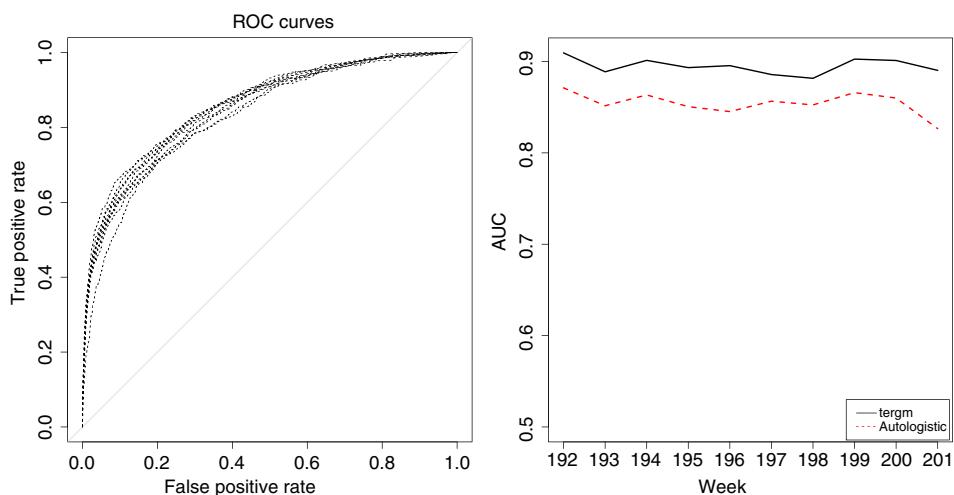


Figure 4. Plots of the 10 operating characteristic curves associated with one-step-ahead out of sample predictions from the autologistic model, and the area under the curves (AUC) for the autologistic model and the temporal ERGM for the NYMEX FTN.

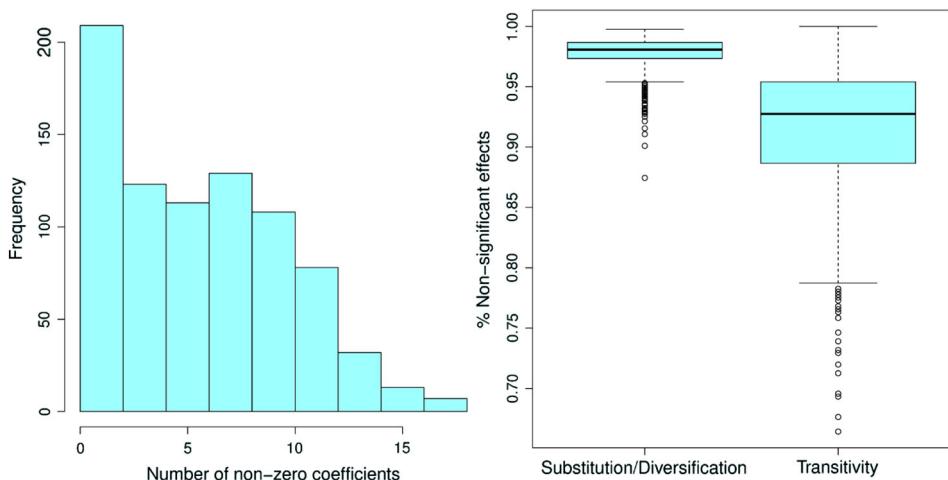


Figure 5. Histogram of the number non-zero coefficients and boxplots of the percentage of non-significant effects of substitution/diversification and transitivity for 812 pairs of traders.

Table 1. Number of effects, and non-presence percentage over 812 pairs of traders.

	# of effects	% non-present
Persistence	6	72.8
Reciprocity	3	88.8
Substitution/diversification	828	0.0
Disintermediation	414	0.4

To understand the overall impact of different trading mechanisms we focus on the 812 pairs that show at least one significant effect and note that a large percentage have no persistence (72.8%) or inter-temporal reciprocity (88.8%) coefficients that are significant (see Table 1). In contrast, each one of these 812 pairs presents at least one substitution/diversification coefficient that might be significant, and the vast majority (99.6%) present at least one potentially significant inter-temporal transitivity effect. This clearly suggests that second-order effects (substitution/diversification and transitivity) are much more important in this FTN than first-order effects (persistence and reciprocity).

To better understand the relative importance of the second-order effects, we also compute the number of non-significant effects of each type for each pair of traders (see Figure 5). Note that, although all 812 pairs present at least one substitution/diversification effect, the number of these effects that might be significant on each pair is relatively small. In contrast, the number of potentially significant coefficients associated with inter-temporal transitivity effects tends to be larger, with a few pairs presenting more than 30% of potentially significant effects. These results suggests that the evolution of this trading network is driven in majority by inter-temporal transitivity effects.

We can contrast these results with those obtained from the tERGM model (see Table 2). Interestingly, all coefficients appear to be significant in this case. Hence, unlike our model, the tERGM suggests that trading is 'sticky' (both momentum effects are significant). This is likely driven by the fact that in the default specification of the tERGM the momentum terms are the only ones that capture the temporal evolution of the network.

Table 2. Point estimates and 95% confidence intervals for the coefficients of tERGM model in the NYMEX FTN data.

Parameter	Estimate
Number of edges	-7.542
Number of paths	0.174
Reciprocity	1.134
Transitive Triplets	-0.043
Cyclic triplets	0.074
Sqrt of in-degree for receiver	1.096
Sqrt of out-degree for receiver	-0.368
Sqrt of out-degree for emitter	0.751
Momentum	0.771
Momentum of reciprocal	0.720

6. Discussion

We introduced a novel statistical model for the analysis of FTNs and applied it to study the NYMEX natural gas futures market between January 2005 and December 2008. Our analysis shows that diversification and substitution effects rather than persistence tend to dominate this market's microstructure.

Our approach focuses on L_1 penalties mainly because of computational expediency. However, alternatives such as the adaptive Lasso [35] or smoothly clipped absolute deviation [8] penalties can potentially improve variable selection. We would also like to explore in the future fully Bayesian implementation using spike-and-slab priors (e.g. see [16]) instead of convex non-differentiable penalties. However, computation for this type of models (particularly for networks with a large number of nodes) is challenging. Furthermore, we note that the model can be easily extended to undirected network by considering a reduced set of predictors, and to weighted networks by replacing the multinomial likelihood with an appropriate member of the exponential family. Similarly, we could extend to model to consider higher order autoregressive processes, but the number of parameters grows dramatically in that case. One way to deal with this issue is to focus on a smaller number of hyperparameters or drop interactions from the model.

In addition to the componentwise maximization algorithm described in Section 4, we also investigated the use of a split-Bregman algorithm [11] and a stochastic gradient algorithm [25], but found both algorithms to have suboptimal performance in this problem. However, for problems with very large T (e.g. high frequency FTNs in highly liquid markets), a stochastic gradient descent approach might be appropriate, and its implementation is relatively straightforward.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This research was partially supported by NSF/DMS award number [1441433].

References

- [1] L. Adamic, C. Brunetti, J.H. Harris, and A.A. Kirilenko, *Trading networks*, Technical report, MIT Sloan School of Management, 2010.

- [2] E. Airoldi, D.M. Blei, S.E. Fienberg, and E.P. Xing, *Mixed membership stochastic blockmodels*, *J. Mach. Learn. Res.* 9 (2008), pp. 1981–2014.
- [3] D.L. Banks and K.M. Carley, *Models for network evolution*, *J. Math. Sociol.* 21 (1996), pp. 173–196.
- [4] B. Betancourt, A. Rodríguez, and N. Boyd, *Modelling and prediction of financial trading networks: An application to the nymex natural gas futures market*, Technical report, University of California – Santa Cruz, 2015.
- [5] P.J. Bickel, Y.A. Ritov, and A.B. Tsybakov, *Simultaneous analysis of lasso and dantzig selector*, *Ann. Statist.* 37 (2009), pp. 1705–1732.
- [6] C. Bliss, M. Frank, C. Danforth, and P.S. Dodds, *An evolutionary algorithm approach to link prediction in dynamic social networks*, *J. Comput. Sci.* 5 (2014), pp. 750–764.
- [7] S.J. Cranmer and B.A. Desmarais, *Inferential network analysis with exponential random graph models*, *Polit. Anal.* 19 (2011), pp. 66–86.
- [8] J. Fan and R. Li, *Variable selection via nonconcave penalized likelihood and its oracle properties*, *J. Amer. Statist. Assoc.* 96 (2001), pp. 1348–1360.
- [9] J. Friedman, T. Hastie, and R. Tibshirani, *Regularization paths for generalized linear models via coordinate descent*, *J. Statist. Softw.* 33 (2010), 1–22.
- [10] A. Goldenberg, A.X. Zheng, S.E. Fienberg, and E.M. Airoldi, *A survey of statistical network models*, *Found. Trends Mach. Learn.* 2 (2009), pp. 129–233.
- [11] T. Goldstein and S. Osher, *The split Bregman method for l_1 -regularized problems*, *SIAM. J. Imaging Sci.* 2 (2009), pp. 323–343.
- [12] S. Hanneke, W. Fu, and E.P. Xing, *Discrete temporal models of social networks*, *Electron. J. Statist.* 4 (2010), pp. 585–605.
- [13] P.D. Hoff, A.E. Raftery, and M.S. Handcock, *Latent space approaches to social network analysis*, *J. Amer. Statist. Assoc.* 97 (2002), pp. 1090–1098.
- [14] P.W. Holland and K. Leinhardt, *An exponential family of probability distributions for directed graphs*, *J. Amer. Statist. Assoc.* 76 (1981), pp. 33–50.
- [15] Z. Huang and D.K.J. Lin, *The time-series link prediction problem with applications in communication surveillance*, *INFORMS J. Comput.* 21 (2009), pp. 286–303.
- [16] H. Ishwaran and J.S. Rao, *Spike and slab variable selection: Frequentist and Bayesian strategies*, *Ann. Statist.* 33 (2005), pp. 730–773.
- [17] S. Kakade, O. Shamir, K. Sindharan, and A. Tewari, *Learning exponential families in high-dimensions: Strong convexity and sparsity*, In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (AISTATS-10)*, Sardinia, Italy, Y.W. Teh and D.M. Titterington, eds., Vol. 9, 2010, pp. 381–388.
- [18] E.D. Kolaczyk, *Statistical Analysis of Network Models*, Springer, New York, NY, 2009.
- [19] B. Krishnapuram and A. Hartemink, *Sparse multinomial logistic regression: Fast algorithms and generalization bounds*, *IEEE. Trans. Pattern Anal. Mach. Intell.* 27 (2005), pp. 957–968.
- [20] J.D. Lee, Y. Sun, and J.E. Taylor, *On model selection consistency of regularized m -estimators*, *Electron. J. Stat* 9 (2014), pp. 608–642.
- [21] P. Leifeld, S. Cranmer, and B. Desmarais, *xergm: Extensions of Exponential Random Graph Models*, R package, 2014.
- [22] M.Y. Park and T. Hastie, *L_1 -regularization path algorithm for generalized linear models*, *J. R. Stat. Soc. Ser. B* 69 (2007), pp. 659–677.
- [23] P. Sarkar and A.W. Moore, *Dynamic social network analysis using latent space models*, *SIGKDD Explorations: Special Edition on Link Mining* 7 (2005), pp. 31–40.
- [24] D.K. Sewell and Y. Chen, *Analysis of the formation of the structure of social networks by using latent space models for ranked dynamic networks*, *J. Roy. Statist. Soc. Ser. C* 64 (2015), pp. 611–633.
- [25] S. Shalev-Shwartz and A. Tewari, *Stochastic methods for l_1 -regularized loss minimization*, *J. Mach. Learn. Res.* 12 (2011), pp. 1865–1892.
- [26] T.A.B. Snijders, C. Steglich, and G. van de Bunt, *Introduction to actor-based models for network dynamics*, *Soc. Netw.* 32 (2010), pp. 44–60.

- [27] R. Tibshirani, *Regression shrinkage and selection via the lasso*, J. R. Stat. Soc. Ser. B 58 (1996), pp. 267–288.
- [28] R.J. Tibshirani and J. Taylor, *Degrees of freedom in lasso problems*, Ann. Statist. 40 (2012), pp. 1198–1232.
- [29] D. Vidaurre, C. Bielza, and P. Larrañaga, *A survey of l_1 regression*, Int. Stat. Rev. 81 (2013), pp. 361–387.
- [30] M.J. Wainwright, *Sharp thresholds for high-dimensional and noisy sparsity recovery using l_1 -constrained quadratic programming*, IEEE Trans. Inform. Theory 55 (2009), pp. 2183–2202.
- [31] E.P. Xing, W. Fu, and L. Song, *A state-space mixed membership blockmodel for dynamic network tomography*, Annals of Applied Statistics 4 (2010), pp. 535–566.
- [32] P. Zhao and B. Yu, *On model selection consistency of lasso*, Journal of Machine Learning Research 7 (2006), pp. 2541–2563.
- [33] Y. Zheng and J. Zhu, *Markov chain monte carlo for a spatial-temporal autologistic regression model*, Journal of Computational and Graphical Statistics 17 (2008), pp. 123–137.
- [34] J. Zhu, Y. Zheng, A.L. Carroll, and B.H. Aukema, *Autologistic regression analysis of spatial-temporal binary data via monte carlo maximum likelihood*, J. Agric. Biol. Environ. Stat. 13 (2008), pp. 84–98.
- [35] H. Zou, *The adaptive lasso and its oracle properties*, J. Am. Stat. Assoc. 101 (2006), pp. 1418–1429.
- [36] H. Zou, T. Hastie, and R. Tibshirani, *On the ‘degrees of freedom’ of the lasso*, Ann. Stat. 35 (2007), pp. 2173–2192.