DECOMPOSING NETWORK DYNAMICS: SOCIAL INFLUENCE REGRESSION

SHAHRYAR MINHAS AND PETER D. HOFF

Abstract. Understanding network influence and its determinants are key challenges in political science and network analysis. Traditional latent variable models position actors within a social space based on network dependencies but often do not elucidate the underlying factors driving these interactions. To overcome this limitation, we propose the Social Influence Regression (SIR) model, an extension of vector autoregression tailored for relational data that incorporates exogenous covariates into the estimation of influence patterns. The SIR model captures influence dynamics via a pair of $n \times n$ matrices that quantify how the actions of one actor affect the future actions of another. This framework not only provides a statistical mechanism for explaining actor influence based on observable traits but also improves computational efficiency through an iterative block coordinate descent method. We showcase the SIR model's capabilities by applying it to monthly conflict events between countries, using data from the Integrated Crisis Early Warning System (ICEWS). Our findings demonstrate the SIR model's ability to elucidate complex influence patterns within networks by linking them to specific covariates. This paper's main contributions are: (1) introducing a model that explains third-order dependencies through exogenous covariates and (2) offering an efficient estimation approach that scales effectively with large, complex networks.

Date: July 24, 2025.

Appendix

Visualization of convergence for direct (blue), sender influence (green), and receiver influence (red) parameters.

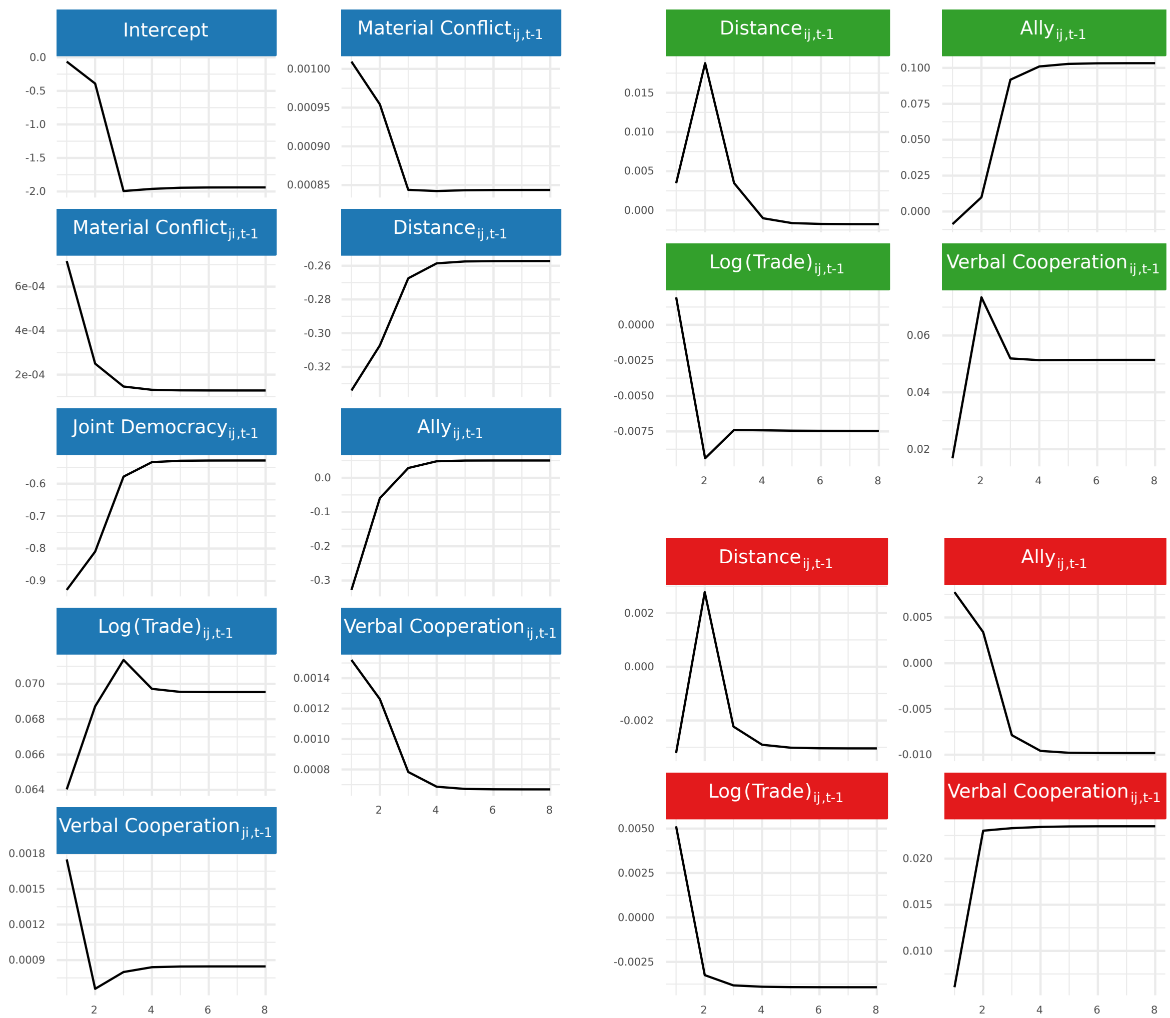

Figure A1. Convergence diagnostics for the social influence regression model on material conflict.

Convergence.

Influence Dynamics. Visualization of influence effects for select time points from dynamic social influence regression model.

Figure A2. Influence relationships

Predictions from SIR. To illustrate how predictions are generated in the SIR model, consider the Poisson case, where

$$y_{i,j,t} \sim \text{Poisson}(\mu_{i,j,t}) \quad \text{and} \quad \mu_{i,j,t} = \exp\Big(\theta^\top z_{i,j,t} + \alpha^\top \tilde{X}_{i,j,t}\, \beta\Big).$$

In this formulation, $z_{i,j,t}$ is a vector of exogenous covariates (e.g., alliance status, trade volume, distance), and $\tilde{X}_{i,j,t}$ typically incorporates terms capturing network dependence, such as a logged lag of the dependent variable ($\tilde{x}_{i,j,t} = \log(y_{i,j,t-1} + 1)$). The parameters θ, α, and β are estimated through the iterative block coordinate descent method described in the manuscript.

Once $\widehat{\theta}$, $\widehat{\alpha}$, and $\widehat{\beta}$ have been obtained, the predicted mean count for any dyad (i, j) at time t is given by

$$\widehat{\mu}_{i,j,t} = \exp\Big(\widehat{\theta}^\top z_{i,j,t} + \widehat{\alpha}^\top \tilde{X}_{i,j,t}\, \widehat{\beta}\Big).$$

These predicted values can then be used for several purposes:

- **In-sample Fitted Values.** Substituting the observed covariates and lagged values into the formula for $\widehat{\mu}_{i,j,t}$ yields fitted counts for the same time period used in model estimation. Comparing $\widehat{\mu}_{i,j,t}$ to $y_{i,j,t}$ helps diagnose how well the model captures the patterns observed in the training data.
- **Out-of-sample Forecasts.** For future time points $t = T + 1, T + 2, \ldots$, one can hold out the corresponding $y_{i,j,t}$ values, estimate the model on $\{1, \ldots, T\}$, and then form predictions by plugging in $z_{i,j,T+1}$ and $\tilde{X}_{i,j,T+1}$. If additional lags are needed, one can either use recently observed values (if available) or replace them with predicted counts from the previous step in a multi-step forecasting procedure. This is what is done in the model evaluation portion of the manuscript.
- **Counterfactual Scenarios.** Since the linear predictor is $\widehat{\theta}^\top z_{i,j,t} + \widehat{\alpha}^\top \tilde{X}_{i,j,t}\, \widehat{\beta}$, one can also modify elements of $z_{i,j,t}$ (e.g., turn an alliance indicator on or off, adjust trade volumes) while keeping all other factors the same. These "what if" scenarios highlight how changes in exogenous covariates shift the predicted count, providing an interpretable gauge of each covariate's substantive impact.

Interpreting Exogenous Covariates. To illustrate how exogenous covariates $z_{i,j,t}$ affect the outcome in the Poisson SIR model, recall that

$$\mu_{i,j,t} \;=\; \exp\Bigl(\theta^\top z_{i,j,t} \;+\; \alpha^\top \tilde{X}_{i,j,t}\,\beta\Bigr).$$

Here, $\mu_{i,j,t}$ represents the expected count for dyad (i,j) at time t. In this section, we focus on the contribution of $\theta^\top z_{i,j,t}$, which captures the direct effect of exogenous dyadic covariates on the log of the Poisson mean.

1. Continuous Covariates. If z_k is a continuous variable in $z_{i,j,t}$, then θ_k indicates how a one-unit increase in z_k shifts $\log(\mu_{i,j,t})$. More precisely, holding all other elements of $z_{i,j,t}$ constant, we have:

$$\mu_{i,j,t} \;\longrightarrow\; \mu_{i,j,t}\cdot e^{\theta_k} \quad \text{as } z_k \text{ increases by one unit.}$$

Thus, if $\theta_k > 0$, increasing z_k will multiply the expected count by e^{θ_k}. Conversely, if $\theta_k < 0$, raising z_k will decrease the expected count by a factor of $e^{\theta_k}\,(< 1)$. This multiplicative interpretation follows directly from the log-link used in the Poisson model.

2. Binary Covariates. If z_k is a binary indicator taking values 0 or 1 (e.g., an alliance indicator), then its effect on the predicted count can be summarized via the difference or ratio of $\mu_{i,j,t}$ under the two states. Specifically, let

$$\mu^{(1)}_{i,j,t} \;=\; \exp\Bigl(\theta^\top z^{(z_k=1)}_{i,j,t} \;+\; \alpha^\top \tilde{X}_{i,j,t}\,\beta\Bigr) \quad \text{and} \quad \mu^{(0)}_{i,j,t} \;=\; \exp\Bigl(\theta^\top z^{(z_k=0)}_{i,j,t} \;+\; \alpha^\top \tilde{X}_{i,j,t}\,\beta\Bigr).$$

The difference $\mu^{(1)}_{i,j,t} - \mu^{(0)}_{i,j,t}$ measures the absolute change in the expected count when z_k moves from 0 to 1, while the ratio $\mu^{(1)}_{i,j,t}/\mu^{(0)}_{i,j,t}$ shows the relative (multiplicative) shift. In many applications, the ratio is more directly interpretable: for instance, a ratio of 1.5 indicates a 50% increase in the expected number of events due to the presence of the condition encoded by z_k.

3. Putting It All Together. Because the full linear predictor is

$$\eta_{i,j,t} \;=\; \theta^\top z_{i,j,t} \;+\; \alpha^\top \tilde{X}_{i,j,t}\,\beta,$$

the exogenous covariates $z_{i,j,t}$ work alongside the terms $\alpha^\top \tilde{X}_{i,j,t}\,\beta$, which encode past network dynamics (e.g., lagged interactions). In practice, researchers often explore changes in $z_{i,j,t}$ while holding $\tilde{X}_{i,j,t}$ fixed at observed or typical levels, thereby isolating the direct effect of an exogenous factor. For instance:

- *One-unit shift in a continuous variable.* Examine how $\mu_{i,j,t}$ is multiplied by e^{θ_k} when z_k moves from $z_{k,0}$ to $z_{k,0} + 1$.
- *Binary switch.* Compare $\mu^{(1)}_{i,j,t}$ and $\mu^{(0)}_{i,j,t}$ to quantify how turning a 0/1 indicator on (e.g., forming an alliance) modifies the expected count of events.
- *Counterfactual scenario.* Set certain components of $z_{i,j,t}$ to alternate values (e.g., higher or lower trade volumes, toggling an alliance indicator) while keeping other aspects of $z_{i,j,t}$ or $\tilde{X}_{i,j,t}$ at their typical levels. This reveals how hypothetical changes in exogenous variables, independent of network history, alter predicted interaction counts.

Through these approaches, one can uncover the substantive roles that external factors (e.g., political ties, economic relationships) play in shaping interactions in the network.

Scoring Rules for Count Data. Scoring rules are penalties $s(y, P)$ introduced with P being the predictive distribution and y the observed value. The goal of researchers interested in prediction is to minimize the expectation of these scores, which is typically calculated by taking the average:

$S = \frac{1}{n} \sum_{i=1}^{n} s(y_i, P_i)$,

where y_i refers to the i^{th} observed count and P_i the i^{th} predictive distribution. A set of proper scoring rules as defined by Czado, Gneiting and Held (2009) are shown in the list below. For each of these rules, $f(y)$ denotes the predictive probability mass function. $\hat{\mu}$ and $\hat{\sigma}$ refer to the mean and standard deviation of the predictive distribution.

- Dawid-Sebastiani score: $s(y, P) = (\frac{y-\hat{\mu}}{\hat{\sigma}})^2 + 2 \times log(\hat{\sigma})$
- Logarithmic score: $s(y, P) = -log(f(y))$
- Brier score: $s(y, P) = -2f(y) + \sum_k f^2(k)$
- Spherical score: $s(y, P) = \frac{f(y}{\sqrt{\sum_k f^2(k)}}$

Author Biographies

Shahryar Minhas is an Associate Professor in the Department of Political Science at Michigan State University.

Peter D. Hoff is a James B. Duke Distinguished Professor of Statistical Science at Duke University.

Shahryar Minhas: Department of Political Science

Current address: Michigan State University

Email address, Corresponding author: `minhassh@msu.edu`

Peter D. Hoff: Department of Statistics

Current address: Duke University

Email address: `peter.hoff@duke.edu`