

This document is the accepted manuscript version of the following article:
Cranmer, S. J., Leifeld, P., McClurg, S. D., & Rolfe, M. (2017).
Navigating the range of statistical tools for inferential network
analysis. *American Journal of Political Science*, 61(1), 237-251. <https://doi.org/10.1111/ajps.12263>

Navigating the Range of Statistical Tools for Inferential Network Analysis*

Skyler J. Cranmer[†]

, Philip Leifeld

, Scott D. McClurg

, and Meredith Rolfe

University of Bern, Institute of Political Science

Southern Illinois University, Department of Political Science

University of Massachusetts–Amherst, Department of Political Science

The last decade has seen substantial advances in statistical techniques for the analysis of network data, and a major increase in the frequency with which these tools are used. These techniques are designed to accomplish the same broad goal, statistically valid inference in the presence of highly interdependent relationships, but important differences remain between them. We review three approaches commonly used for inferential network analysis—the Quadratic Assignment Procedure, Exponential Random Graph Model, and Latent Space Network Model—highlighting the strengths and weaknesses of the techniques relative to one another. An illustrative example using climate change policy network data shows that all three network models outperform standard logit estimates on multiple criteria. This paper introduces political scientists to a class of network techniques beyond simple descriptive measures of network structure, and helps researchers choose which model to use in their own research.

*SJC's work on this project was supported by the National Science Foundation (SES-1357622, SES-1461493, and SES-1514750) and the Alexander von Humboldt Foundation's fellowship for experienced researchers.

[†]Please direct correspondence to cranmer.12@osu.edu.

Replication Materials: The data, code, and any additional materials required to replicate all analyses in this article are available on the American Journal of Political Science Dataverse within the Harvard Dataverse Network, at: <http://dx.doi.org/10.7910/DVN/2XP8YF>

Statistical methods usually require observational independence, or the assumption that actors do not influence each other's outcomes. While such an assumption may apply to randomized surveys or lab experiments, it is less believable in social and political systems where the interactions between actors are an integral part of the process of interest. The implications of *interdependencies* have been recognized in the cases of time series and spatial data, but political scientists are less familiar with the methods used to model network-based interdependence. Crucially, the development of sophisticated statistical models for studying interdependence is not “merely” a statistical advance, but opens the door for new questions and theories to be posed and tested. Our goal is to introduce and compare three methods commonly used for statistical inference with network data, and to demonstrate the added value of specialized network models compared to a more traditional logit-based estimation approach.

Here, we focus on three techniques designed to model relational outcomes: the Quadratic Assignment Procedure, Exponential Random Graph Models and Latent Space Network Models. In other words, these are models that describe a network of political interactions, where the outcome of interest is a political relationship (e.g., wars, legislative collaboration etc.). Such an understanding of networks—that is to say, one focused on explaining interdependence in relationships—is of interest in many corners of the discipline. For example, networks are formed by terror groups (??), mutual protection treaties (?), and trade flows (?). Institutional networks are relevant to diaspora politics (?), the impact of nongovernmental organizations on policy (?), and cooperation and conflict in legislatures (?). Networks are also part of understanding political behavior, especially where we are interested in questions involving integration, cooperation, and conflict (??).

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1111/ajps.12263

This article is protected by copyright. All rights reserved.

Political Science and the “Problem” of Interdependence

The study of interdependence is part of the core mission of the social sciences, if not *the* core mission (?). It is impossible to describe political and social activity without describing how actors (individuals, groups, states etc.) interact with, react to, adapt to, and influence one another in complex and dynamic ways. One can view interdependence as either a threat to the validity of statistical analysis, or as fundamental to a theory of political processes. These two perspectives are radically different in substance, yet both imply the need to model interdependencies empirically. While we hold that the latter perspective is more compelling, political scientists have generally viewed interdependence as a threat.

Most statistical techniques, including regression models, depend critically on the presence of conditional independence of observations for their validity. Consider a generalized linear model estimated by maximum likelihood (e. g., logit, Poisson, or Gaussian linear model). The joint likelihood—that is, the object to be maximized in order to recover the parameter estimates—is computed as:

$$L(\theta|\mathbf{y}, \mathbf{X}) = \prod_{i=1}^n f(y_i|\mathbf{x}_i, \theta), \quad (1)$$

where θ are the parameters, \mathbf{y} is the vector of outcomes, \mathbf{X} is the matrix of covariates, n is the number of observations in \mathbf{y} , and y_i is the outcome for a given observation i (?). In other words, the likelihood for all the data in the model is the product over the individual functions of each observation i . Because each observation is entirely independent of every other observation, conditional only on the covariates included in the \mathbf{X} matrix, the axioms of probability dictate that the product over these individual values is a valid joint likelihood for all the data.

Critically, this joint likelihood will *only* be valid if the observations are conditionally independent. If the conditional independence assumption is broken, by omission of a relevant covariate (as in the case of omitted variable bias) or the omission of relevant endogenous (network) dependencies, Equation (1) will *not* be a joint likelihood. To be clear, nothing will block the computation of *a* quantity,¹ but that quantity is not the joint likelihood

¹Statistical software has no means of knowing whether the independence assumption

and has no known statistical properties. The bias introduced by violation of the independence assumption is arbitrary and can range from trivial to catastrophic, contingent on the extent of the un-modeled dependencies and the sensitivity of the model to them.

Substantively, we must ask ourselves whether the observations under consideration influence each other through means not accounted for by the covariates. ? illustrate the limits of the independence assumption in an analysis of World War II. The assumption of conditional independence in a traditional logit model of conflict implies that the British declaration of war on Germany in 1939 was *entirely* independent of all factors other than, say, regime type similarity and the ratio of relative capabilities between the two states. Such a claim is comical, as Britain’s declaration of war was unambiguously a response to Germany’s invasion of Poland, but standard regression models cannot account for such dependencies.

Several recent studies have shown that, when network dependencies are properly accounted for, longstanding substantive findings may be called into question. Perhaps the most persuasive example of this involves the “democratic peace,” being the idea that democracies tend not to fight one another (see ? and ? for canonical works). The effect of the democratic peace has been largely established via logistic regression models, which assume conditional independence of observations. Recent works by ? and ? have shown that joint democracy has no statistically discernible effect on conflict once network dependencies are accounted for, thus suggesting that network structure in which conflicts occur, rather than something about the nature of joint democracy, accounts for the paucity of conflicts between democratic states. What is more, failure to *properly* model network dependencies can result in the same sort of bias. ? were interested in the network concept of structural equivalence, but included this effect as a predictor in a logistic regression. A replication of this analysis by ? using a temporal ERGM uncovered substantial bias in the results, including a new lack of statistical significance for their operationalization of the democratic peace.

is violated and no errors will be produced by even wild violations of this assumption.

Models for Statistical Inference on Relational Data

While many network analysis techniques are descriptive in nature, we review and compare three methods commonly used for inferential network analysis: the Quadratic Assignment Procedure, Latent Space Model, and Exponential Random Graph Model. We focus on cross-sectional models where binary network ties (called *edges*) between actors (called *vertices*) are the outcome of interest (e. g., relational data where the outcome of interest is the presence or absence of an edge), although these models can often be extended to accommodate different network types (e. g., longitudinally observed networks and networks with weighted edges).

The Quadratic Assignment Procedure (QAP)

The Quadratic Assignment Procedure (QAP) is a non-parametric test for the significance of an association between two matrices with complex dependencies (?). QAP is not a statistical model but an add-on for standard regression models that provides an unbiased test of association in spite of potential dependencies. While the original QAP (?) was designed for bivariate associations, ?? introduced Multiple Regression QAP (MRQAP).

Assuming that linear models for square matrices take the form

$$N = \beta X + \gamma Z + E, \quad (2)$$

where N , X , Z and E are $n \times n$ matrices (with n denoting the number of vertices in the network), X and Z are predictors, N is the outcome network, E is the error term, and β and γ are coefficients to be estimated, the uncertainty measures of β and γ are generally biased because the observations within N are partly dependent on each other.

The intuition behind the QAP is that permutations of N , where the order of the rows and columns are reshuffled simultaneously in the same way, preserve the dependency structure across observations within N , but remove the dependencies in the associations between N and X as well as N and Z . For example, if X denotes a matrix containing legislators' common memberships in committees and N denotes legislators' information exchange patterns, then a simultaneous permutation of the rows and columns of N (in the

same way) would still preserve the information on who exchanges information with whom, but it is no longer possible to relate the common committee memberships of a specific actor pair (ij) to the information exchange pattern of the same dyad. While there is likely an association between X and the original, non-permuted N , there is not expected to be an association between X and the permuted N .

The QAP uses this permutation procedure to create an empirical sampling distribution reflecting a new null hypothesis of no association that correctly takes into account the correlation between observations. The procedure repeats the permutations a large number of times and recomputes the association measure, for example the regression coefficient for X in the MRQAP case, at each iteration. The resulting distribution corresponds to the distribution under the null hypothesis that no association between X and N exists, but correctly takes into account the dependencies within N . Finally, the original coefficient for X is compared to the new distribution to see how many elements of the new distribution are less extreme than the coefficient, in order to obtain a dependency-corrected (and thus bias-corrected) p value for the coefficient. Essentially, this approach treats dependencies among the observations as a nuisance that needs to be corrected for rather than as an interesting feature that can be modeled. This makes the approach easier to learn and apply, but offers no modeling choices if one is substantively interested in the structure of the dependencies.

? generalizes MRQAP to arbitrary outcome distributions, for example the common case of binary matrices, by applying the QAP test to coefficients from generalized linear models. However, subsequent research on the original QAP test finds that the p values produced by the original permutation approach are prone to bias under certain conditions (?): (combinations of) non-pivotal association tests (such as partial regression coefficients, see ? and ?), collinearity (the presence of a third variable Z in the model, which has a correlation both with N and the other predictor X , see ?), small numbers of vertices (?), skewness of the edge weights in the outcome network (?), and high levels of autocorrelation between observations of the same vertices (?) lead to biased uncertainty measures. The current state of the art are “Freedman–Lane Semi-Partialing” (?) and “Double-Semi-Partialing” (?), both of which resolve most of the problems mentioned above when the general linear model is used (?). Freedman–Lane Semi-Partialing and Dekker’s Double-Semi-Partialing are termed “residual methods” (as opposed to “raw methods”, which permute the raw N matrix) because they permute the resid-

uals in order to partial out the effect of the third variable Z , which partly explain X and N . The Freedman–Lane Semi-Partialing method first estimates the effect of Z on N alone (without including the model term of interest X), then permutes the residuals of that model to simulate a distribution of new N matrices, and finally recomputes the full model for all simulated N in order to obtain a sampling distribution, thereby partialing out the effect of Z from N . Dekker’s Double-Semi-Partialing is similar, but partials out the effect of the third variable Z from the effect of the other predictors X , rather than the response matrix N . ? show that these two approaches are largely unbiased under most circumstances in conjunction with a linear model; only high levels of collinearity between X and Z in combination with highly skewed edge values in the dependent network lead to an increased number of type I errors.

It is yet unclear how this result would extend to binary network matrices with few edges (and thus a skewed distribution) when collinearity is present—despite numerous published applications to binary networks in conjunction with logistic regression (e. g., ?). Beside this caveat, standard errors cannot be recovered when the QAP approach is used; only the coefficient and its p value are available. Moreover, QAP-corrected models are subject to the same limits as their uncorrected counterparts, for example they are still sensitive to omitted variable bias.

However, the MRQAP has some benefits that make it worth considering. Critically, it is the most accessible of the methods treated here. It is well implemented² and easy to interpret as MRQAP results can be interpreted like other regression. MRQAP can also produce reasonable results under conditions that might make it difficult to obtain estimates using other methods (such as a dense network of relationships). Finally, the MRQAP does not need a theoretical model to “deal with” network dependencies. While this reduces the likelihood that the results are contingent on a theoretically misspecified model, that also implies that there are no opportunities to learn about those dependencies.

Exponential Random Graph Models (ERGMs)

The exponential random graph model, or ERGM, is a powerful model that has seen increased usage in recent years. First proposed by ?, the ERGM

²Implementations include `Ucinet` (?), `Stata`, and the `sna` package (?) in R.

is a direct operationalization of the joint probability density from which the networks are thought to be generated, with minimal modeling assumptions. For intuition, consider that the ERGM finds its parameters by maximizing the probability of the observed network over the networks with the same number of vertices that could have been observed. This is conditional on a set of network statistics that can include vertex- and edge-level exogenous variables and endogenous dependencies; a substantial increase in the scope of modelable dependencies without any independence assumptions. This produces a single, multivariate, distribution from which the network of interest is the most likely to be drawn.

The ERGM is structured as follows:

$$\mathcal{P}(N, \boldsymbol{\theta}) = \frac{\exp\{\boldsymbol{\theta}'\mathbf{h}(N)\}}{\sum_{N^* \in \mathcal{N}} \exp\{\boldsymbol{\theta}'\mathbf{h}(N^*)\}} \quad (3)$$

$\mathbf{h}(N)$ is a vector of statistics computed on the network N with the same number of elements as $\boldsymbol{\theta}$, and \mathcal{N} refers to the set of all possible permutations of the network N —from entirely empty to entirely complete—with the same number of vertices (?). As such, the ERGM literally computes the probability of the network we observed over the networks we could have observed.

The ERGM requires two assumptions. First, one must assume that there is equal probability of observing any two networks with the same values of the statistics included in the specification. This is functionally the same as assuming that a model is completely and correctly specified. Consider two networks that are identical with respect to the values of the statistics computed in $\mathbf{h}(N)$, for example the number of closed triads. If one of these networks is more likely than the other, something not included in the model is causing that imbalance. By implication the model would be incompletely or incorrectly specified. The second assumption is that the observed network exhibits the average value of those statistics over the networks that could have been observed. This assumption is necessary to identify the parameters, but is not different in practice from the assumption made in regression models that the average relationships in a dataset are representative of the population (?).

All effects in the ERGM are included as statistics computed on the network in the vector $\mathbf{h}(N)$. Exogenous covariates, both at the edge and vertex

levels, are included in the ERGM through statistics computed as

$$h_X(N, X) = \sum_{i \neq j} X_{ij} N_{ij}. \quad (4)$$

This statistic accepts a relational matrix X with the same dimensions as the outcome network N and simply sums over a dyad-wise product. So, for the case where X captures a relational covariate measured at the ij dyad level, the most natural type of covariate for this model, one is simply summing over the values of X_{ij} for which the corresponding edge in the outcome network N_{ij} exists. Even though relational covariates are included most naturally, vertex level covariates are also included easily: one need only to “scale up” the vertex level covariate Z_i such that $X_{ij} = Z_i$.

The ERGM is also designed to accommodate a great variety of endogenous dependencies, also as statistics in the $\mathbf{h}(N)$ vector. However, since the endogenous effects will capture different forms of interdependence, they must be specified differently and no single formula exists for articulating them as in Equation (4). Consider two examples of endogenous dependencies includable in the ERGM. First, reciprocity. This effect is endogenous because it occurs only within the outcome network (e. g., i considers j a friend and j considers i a friend as well). Reciprocity is illustrated in the left cell of Figure 1 and may be written as

$$h_R(N) = \sum_{i < j} N_{ij} N_{ji}, \quad (5)$$

Where we see that this statistic counts the number of reciprocal relationships in the network. Second, a slightly more complicated example is the transitivity effect. Transitivity is specified, in a directed network, as

$$h_T(N) = \sum_{i, j \neq i, k} N_{ij} N_{ik} N_{jk}, \quad (6)$$

and illustrated in the right cell of Figure 1. Nearly any form of interdependence can be included in the ERGM specification. The major restriction on the specification and inclusion of such effects is that they must be articulable as sums of sub-network products. See ? for an extensive treatment of the dependence statistics includable in an ERGM. The ERGM is also special in that, unlike many statistical models, it represents a complete and

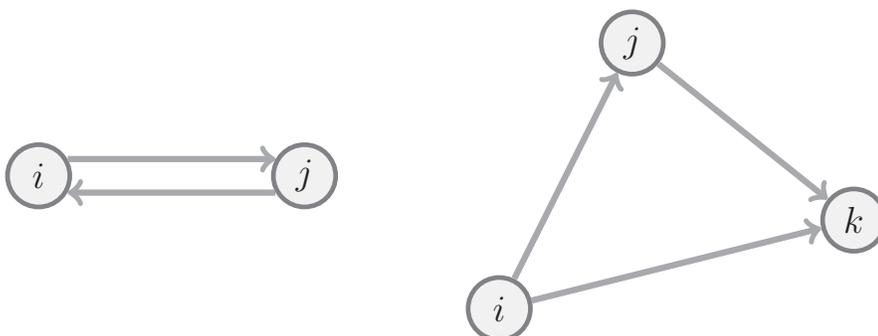


Figure 1: Examples of canonical ERGM statistics. Left panel: reciprocity in a directed network. Right panel, transitivity in a directed network.

proper probability model of the entire network and network data generating process.³

ERGMs are generally estimated in one of two ways. Markov chain Monte Carlo maximum likelihood estimation (MCMC-MLE) is the preferred method for estimating cross-sectional (single network at a single time period) ERGMs (?????). An alternative method of estimation is called maximum pseudolikelihood estimation (MPLE) (?). MPLE is substantially less computationally difficult than MCMC-MLE, but has the problem that confidence measures are downwardly biased. This problem is correctable when using this technique on longitudinally observed networks (?). For details on both of these estimation routines and a direct comparison of the two, see ?.

Yet the ERGM is not without its disadvantages. First, we mentioned that the ability to model specific forms of interdependence is an advantage of the ERGM, and it is, but it is also a disadvantage. Because an ERGM without any endogenous dependencies in its specification reduces to a standard logistic regression model, the researcher *must* model the endogenous dependencies and must also specify them in a complete and correct manner

³While we have discussed the ERGM in its basic form, which is tailored to binary cross-sectional networks, the ERGM extends to longitudinally observed networks through a technique called the Temporal ERGM (TERGM) (??) or a related technique called the stochastic actor oriented model (SAOM) (??) that assumes a more specific updating process between observed networks (?). The ERGM also extends to networks with valued edges (called the Generalized ERGM or GERGM) (?). A detailed treatment of these extensions is beyond the scope of the present discussion, but we do note that both extensions maintain the basic conceptual logic of the ERGM and both are well implemented in the `xergm` package (?).

(otherwise the model will be misspecified). This places an increased burden on the researcher, who may be primarily interested in the role of an exogenous covariate, to accurately model the generative structure of the network. Researchers wishing to avoid this additional modeling duty will find techniques like the QAP and latent space model more appealing than the ERGM because such models condition out dependencies without requiring the researcher to specify them. A second disadvantage of the ERGM is that it can be prone to numerical instability in the estimation process, even when there is a theoretically intuitive specification. In particular, this problem makes the effective estimation of model specifications that fit the data very poorly difficult to achieve. This is perhaps a blessing in disguise as it means that poor models cannot be fit (?), but making sure that one’s model fits the data reasonably well can be challenging, especially for networks in which the density of connections is either very sparse or very dense.

Latent Space Network Models

Latent space models were introduced by ? and operationalize dependence between observations through the notion of a k -dimensional “social space.” The idea is that relations are transitive: if vertices i and j are tied and vertices j and k are tied, i and k are also likely to be tied directly. The positions of vertices in the latent space retain this transitivity in terms of the distances between vertices: vertex pairs with a large (Euclidean or other) distance between each other in the latent space have large path distances in the observed network while proximate vertex pairs in the latent space have small path distances in the observed network. Latent space models are generalized linear models on dyadic observations which control for these dependence structures by conditioning on the distance between the vertices’ latent space positions.

In these models, the observations N_{ij} in network matrix N are conditionally independent given dyadic covariates \mathbf{x}_{ij} , their parameters θ , and the positions z_i and z_j in the latent social space (for details on this exposition, see ?). Thus the probability of observing the network is the product of the individual probabilities for each dyadic observation, given the covariates, their estimates, and the latent space positions of i and j :

$$P(N|Z, X, \theta) = \prod_{i \neq j} P(N_{ij}|z_i, z_j, \mathbf{x}_{ij}, \theta). \quad (7)$$

In other words, the dependencies between different dyads N_{ij} are relegated to the coordinates of individual vertices i and j in the latent space, and these positions are estimated along with the parameters for the covariates. The estimation of positions for all vertices in the network adds a substantial number of parameters to the model and, in this way, a latent space network model may be thought of as an elaborately specified random effects model.

Like in the ERGM and QAP context, the \mathbf{x}_{ij} covariates can include relational information (e.g., do i and j have the same attribute value?) or information on any of the two vertices involved (e.g., an attribute of i irrespective of the attribute value of j or an attribute of j irrespective of the attribute value of i). There can be multiple covariates, and all covariates are saved in matrices conforming to the dimensions $n \times n$ of the outcome network matrix N . This means that the practical data preparation needed to run a latent space model is identical to that needed to run an ERGM.

The dyadic tie probabilities η_{ij} can be expressed as log odds, which results in a logistic regression equation. The probability of each pair of vertices N_{ij} being tied is a function of the distance between i and j in the social space. This distance is most often the Euclidean distance between the positions z_i and z_j over the k dimensions of the latent space:

$$\eta_{ij} = \log \text{odds}(N_{ij} = 1 | z_i, z_j, \mathbf{x}_{i,j}, \alpha, \beta) = \alpha + \beta' \mathbf{x}_{i,j} - |z_i - z_j|. \quad (8)$$

Alternatively, other distance measures between z_i and z_j can be chosen (e.g., $+\frac{z_i z_j}{|z_j|}$ instead of $-|z_i - z_j|$; for an example, see ?). However, the distances must satisfy the triangle inequality,

$$d_{ij} \leq d_{ik} + d_{kj} \quad \forall \{i, j, k\}, \quad (9)$$

meaning that any direct distance between i and j should not be larger than distances over indirect paths between i and j through other vertices k . This condition essentially constitutes the latent social space.

The goal of the estimation process in a latent space model is primarily to find a distance matrix D that satisfies the triangle inequality and thus maps the dependence structure in the network matrix to a matrix of distances where adjacent vertices in the network have smaller distances than non-adjacent vertices, controlling for the covariates. This is done by maximization of

$$\log P(N|\eta) = \eta_{ij} y_{ij} - \log(1 + e^{\eta_{ij}}), \quad (10)$$

where the distance matrix and covariates enter the likelihood via η_{ij} as defined in Equation (8). A matrix of initial distance values must be provided to the estimation routine. It is often sufficient to compute a distance matrix based on N using a measure such as the path distance between any i and j or the Euclidean distance between the rows of N . In more complicated cases, it may be necessary to maximize the log likelihood without covariate parameters first in order to get estimates for D and then use these estimates as initial values for the full optimization of the likelihood including the covariates as defined in Equations (8) and (10). After computing the initial distances, the vertex positions in the latent space Z can be recovered by applying multidimensional scaling (MDS) to the distance matrix. As different MDS solutions with arbitrarily rotated, reflected, or translated coordinates yield the same fit, a Procrustes transformation is applied to the MDS coordinates (for details, see ?).

For illustration, suppose the network of international militarized interstate disputes is to be modeled. First, one must define the relational conflict matrix (N) and the covariate matrices X (e.g., trade, contiguity etc.). Second, one computes a matrix of, e.g., the path distances in the conflict network. Third, one recovers each country's coordinates in a low-dimensional space from the path distance matrix by multidimensional scaling (for a primer, see ??). Fourth, one uses these rough positions as starting values for non-linear maximization of the log likelihood in conjunction with a Euclidean model of space (as defined in Equation 8). This yields estimates for the latent positions Z controlling for the covariates X . The covariate estimates can be interpreted substantively because network dependence has been controlled for via the estimated latent positions, and the positions can be used for description of the network, similar to descriptive uses of techniques like multidimensional scaling (??) when applied to network data.

Compared to the QAP and ERGM, latent space models have a number of advantages and disadvantages. Often cited as an advantage is the fact that latent space models can be applied to outcome networks with binary or valued data of any distributional shape. Yet the same is true for the QAP, and a recent ERGM advance in the form of the Generalized ERGM (GERGM) renders the same true for ERGMs (?).

Latent space models relegate network dependencies to the latent social space and estimate them automatically. This makes latent space models easy to use. Yet the model is less flexible than an ERGM because substantive theory related to the dependencies cannot be tested. In this sense, latent

space models have a similar strength (and weakness) as the QAP. Yet one is required to specify the distance measure and the number of dimensions of the latent social space, which puts latent space models between QAP and ERGMs with regard to the theoretical elaboration and required user skills. The choice of the distance measure and the number of dimensions is rarely guided by theory; it is a tradeoff between model fit and parsimony because more dimensions generally accommodate the dependencies better, but each additional dimension introduces nonlinearly more parameters. For illustration, a latent space model typically has an intercept, several coefficients for covariates, and n parameters per dimension of the latent space. This reveals a potential disadvantage of latent space models vis-à-vis the QAP and ERGM: they explain the same network structure but need many more parameters and are therefore less parsimonious. This is especially problematic in situations where few vertices are present because there may be ultimately more parameters than observations.

Finally, interpretation of the coefficients for exogenous covariates may not always be straightforward: if the latent positions or distances are correlated with exogenous covariates, the coefficients may in fact indicate the *opposite direction* of the actual effect in the presence of the confounding positions or distances. While endogenous model terms may be correlated with exogenous effects in ERGMs as well, these problems are easier to detect and fix in ERGMs.

Weighing these advantages and disadvantages, latent space models are an attractive choice if there are no or few isolates, if parsimony is not of primary importance, if there are enough observations, if the interdependencies are not theoretically interesting, and/or in situations where it makes sense to visualize and interpret the latent positions (e. g., political ideology). There are better models available when the dependencies are substantively interesting or when many isolates or few observations exist.

Illustrative Example: Policy Network Collaboration

To illustrate the strengths and weaknesses of the three models, we apply each to a simple, cross-sectional policy network on relations between 34 political actors in the Swiss climate change mitigation network (?). We re-analyze

her dataset with a logit model and the network methods described above. Our results show that all three network models outperform the standard logit estimates on multiple criteria, with the ERGM performing best overall.⁴

The directed outcome network reflects collaboration among four government agencies, five political parties, six scientific/research organizations, eleven organized interest groups (private sector and business associations), and seven environmental NGOs (??).

Exogenous Dependencies

We follow ? with respect to model specification. See the supporting information for a detailed theoretical discussion. We break down our exogenous specification by theoretical mechanism.

1. **Conflicting policy preferences.** We expect little collaboration between conflictual interests, such as business and environmental groups.
 - (a) **Business vs. NGO** (expected effect: negative). Whether the tie sender is an NGO and the potential receiver is a business association or vice-versa.
 - (b) **Preference dissimilarity** (expected effect: negative). A dissimilarity matrix between the 34 actors over four important policy issues using Manhattan distances.
 - (c) **Opposition/alliance** (expected effect: positive). A measure of *perceived* policy similarity.
2. **Transaction costs.** The literature suggests other mechanisms for collaboration between actors, such as transaction costs associated with acquiring contacts (?).
 - (a) **Joint forum participation** (expected effect: positive). Indicator for whether i and j are members of the same forum.
3. **Political Influence.** Actor i tends to collaborate with j if i deems j influential in the policy process because this is instrumental for achieving policy objectives (?).

⁴This result is robust to multiple specifications of the logistic regression, see the supporting information for further details.

- (a) **Influence attribution** (expected effect: positive). Indicator for whether an organization considers another organization “particularly influential.”
- (b) **Alter’s influence in-degree** (expected effect: positive). Each cell X_{ij} contains the in-degree of column actor j in the influence attribution network.
- (c) **Influence absolute diff.** (expected effect: negative). Absolute difference between the influence in-degree of i and the influence in-degree of j
- (d) **Alter = Government actor** (expected effect: positive). Tests whether state actors are popular collaboration targets.

4. Functional actor role requirements.

- (a) **Ego = Environmental NGO** (expected effect: positive). Indicator for whether ego is an environmental NGO.
- (b) **Same actor type** (expected effect: positive). Indicator for whether i and j are the same type of organization, and thus likely functionally interdependent.
- (c) **Mutuality** (expected effect: positive). If i collaborates with j , j likely collaborates with i (Equation 5 and Figure 1).⁵

Endogenous Dependencies

The QAP and latent space models do not require or allow explicit specification of endogenous dependencies, only the ERGM does that. The downsides of this flexibility are that a theoretical understanding of the endogenous part of the data-generating process is required and the analyst must be familiar with the different ways that dependencies can be modeled. In the ERGM reported below, we include the following endogenous model terms.

1. **Two-paths** (expected effect: negative). This term counts the number of two-paths, configurations of ties ijk where vertices i and k are not directly tied, in the collaboration network. Effectively, the coefficient indicates whether two-paths are more or less prevalent than in a random

⁵Note that this endogenous effect can be included as a covariate in a simple logistic regression model, which is not possible for higher-order dependencies like triadic closure.

graph. We expect a negative coefficient because collaboration often expands to more than two interconnected actors, leading to a tendency of triadic closure rather than open triangles/two-paths.

2. **Edge-wise shared partners (GWESP)** (expected effect: positive). A term that counts the number of shared partners (e. g., transitive indirect ties) for each tied actor pair and adds up the shared partners to the count, but places lower weight on each additional shared partner. This geometric decay is governed by an α parameter. For instance, $\alpha = 0$ would place the same weight on edges with one shared partner and edges with any number of shared partners. The parameter is set by the user, and its choice is usually guided by model fit. A positive GWESP coefficient means that there is a tendency for multiple indirect ties of length two between directly tied actors. Transitive relationships are presumably formed because actors find it less costly to connect to other actors directly if an indirect tie already exists, which makes it a mechanism along the lines of transaction cost politics (?).
3. **Outdegree popularity** (expected effect: positive). This term captures the tendency of i to connect to j the more outgoing ties j has. Intuitively, actors with many ongoing collaborations are attractive collaboration partners for new actors. The term therefore operationalizes clustering of collaborations around some active vertices in the network.
4. **GWegree and GWdegree** (expected effects: positive). These terms capture activity and popularity, respectively. They parametrize the distributions of the number of incoming and outgoing collaborations per vertex using an α decay parameter. Intuitively, larger numbers of incoming (outgoing) ties are geometrically discounted when the indegrees (outdegrees) are summed in the statistic. These terms therefore capture network structures with some highly popular or active agents.

Goodness of Fit Comparison

Before comparing the substantive results of the four models, it is important to assess their respective fit. One of the many advantages of network analysis is that it is generally a simple matter to check the overall fit and performance of a model. Goodness-of-fit assessment, prior to consideration of the substantive results, is important. After all, if the model fits the data poorly, one would

not expect the results to have any bearing on the data generating process and so they would be of no interest to the analyst. We consider that a model should pass at least a basic test of its fit to the data prior to the analyst investing the time to understand that model's substantive results.

Figures 2, 3, 4, and 5 plot several characteristics of model-based simulations against the observed collaboration network for each of the models. This is an easy way to determine whether a model captures network dependencies. For example, simulated networks should have the same distribution of path distances or roughly the same numbers of shared partners per dyad as the observed networks. Failing to capture the dependencies is a form of omitted variable bias. The diagrams show distributions of several typical endogenous network properties that can serve as benchmark criteria for comparing the observed network with 100 simulated networks based on the model and the covariates. If the model captures the dependencies well, the black line, which represents the values of the observed network, should pass approximately through the medians of the boxplots, which represent the distributions of the network statistics in the 100 simulated networks.

The estimated ERGM (Figure 5) is relatively accurate with regard to these auxiliary statistics, with very few exceptions: indegrees of 4, 5, and 16 and outdegrees of 6 are slightly underrepresented in the ERGM, but adjusting the model to accommodate these oddities would likely mean overfitting the data. The logistic regression (Figure 2), in contrast, displays much stronger deviations between simulations and observed network. The edge-wise shared partner, indegree, and outdegree distributions of the simulations and dyads with zero shared partners are not in line with the observed network. The QAP model (Figure 3) and the latent space model (Figure 4) fare between the logit and the ERGM in terms of endogenous fit: in the latent space model, indegree and outdegree, and edge-wise shared partners are more accurate than in the logit model, but not as accurate as in the ERGM; zero dyad-wise shared partners fit well (in contrast to the logit model), but shortest paths of length two and infinity are slightly off the mark. The QAP model has similar fit to the logit; only in-stars fit somewhat better. The conclusion from this goodness-of-fit assessment is that the ERGM is largely unbiased because its endogenous network dependencies are correctly specified, while latent space, QAP and logit models are less accurate, in descending order of quality.

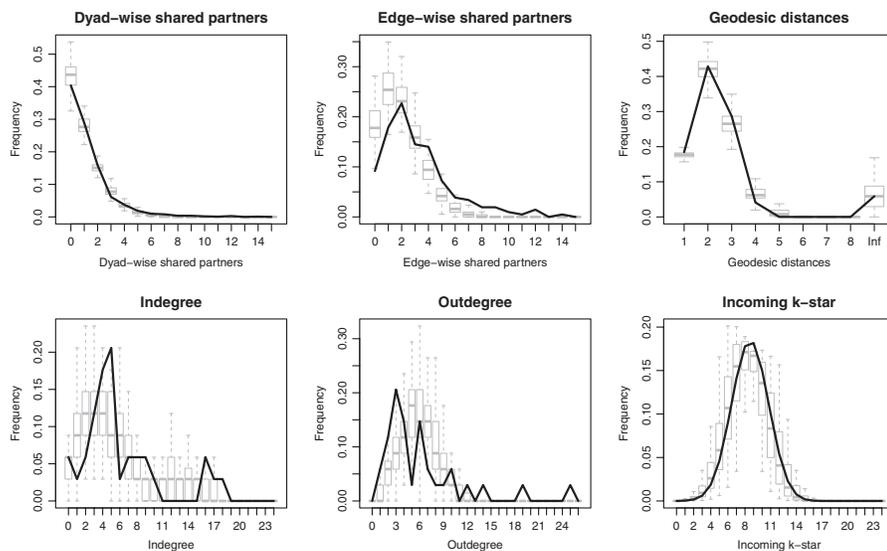


Figure 2: Logit model: goodness-of-fit assessment.

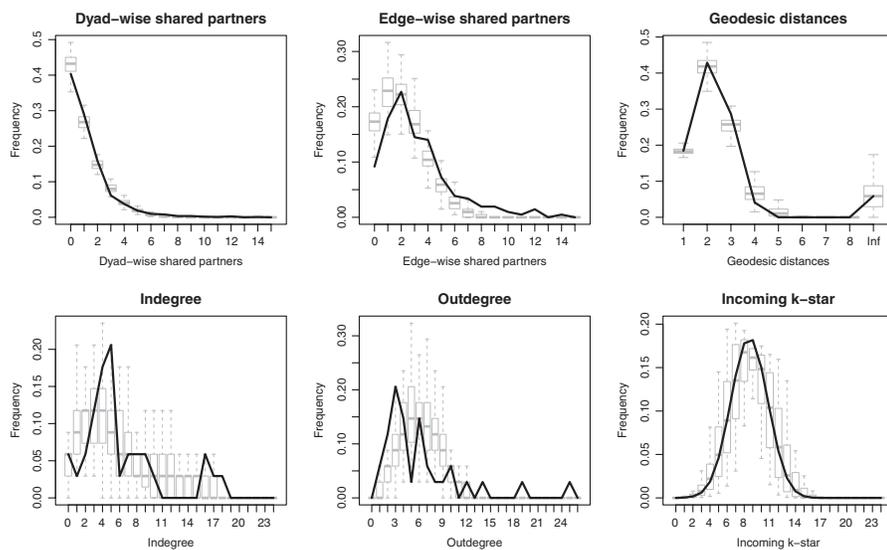


Figure 3: MRQAP model: goodness-of-fit assessment.

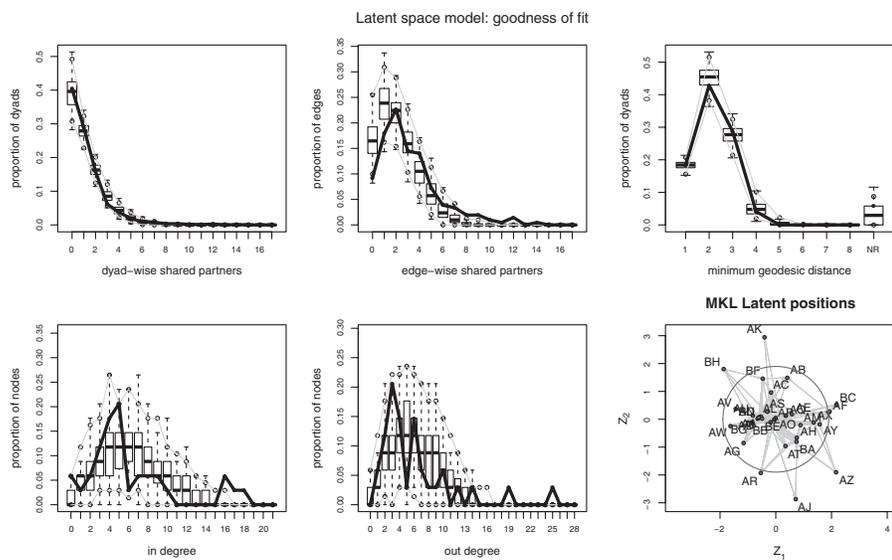


Figure 4: Latent space model: goodness-of-fit assessment and estimated positions of actors.

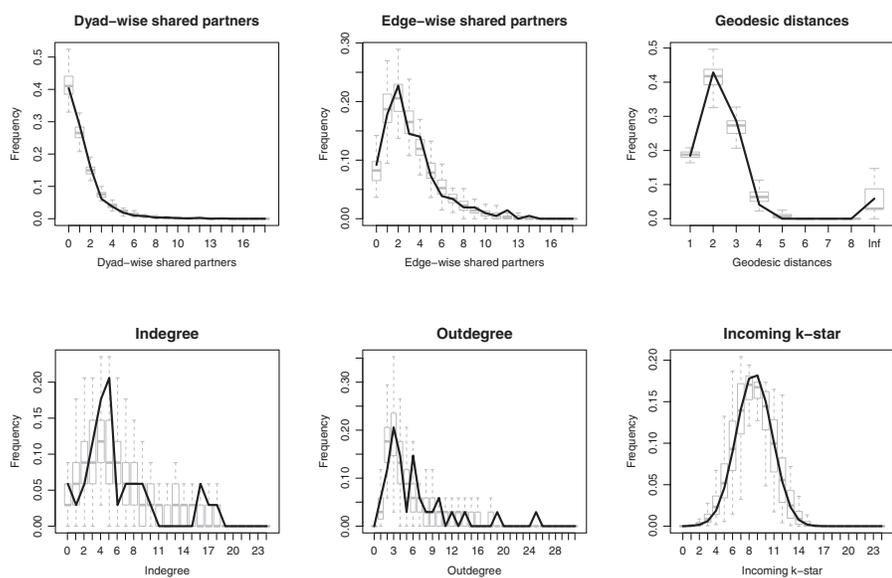


Figure 5: ERGM: goodness-of-fit assessment.

	Logit	MRQAP	LSM	ERGM
Intercept/Edges	-4.44* (0.34)	-4.24*	0.95* [0.10; 1.82]	-12.17* (1.39)
Conflicting policy preferences				
Business vs. NGO	-0.86 (0.46)	-0.87*	-1.37* [-2.40; -0.41]	-1.11* (0.51)
Opposition/alliance	1.21* (0.20)	1.14*	-0.00 [-0.40; 0.40]	1.22* (0.20)
Preference dissimilarity	-0.07 (0.37)	-0.60	-1.76* [-2.60; -0.91]	-0.45 (0.39)
Transaction costs				
Joint forum participation	0.88* (0.27)	0.75*	1.52* [0.86; 2.19]	0.90* (0.28)
Influence				
Influence attribution	1.20* (0.22)	1.29*	0.08 [-0.41; 0.54]	1.00* (0.21)
Alter's influence indegree	0.10* (0.02)	0.11*	0.01 [-0.03; 0.04]	0.21* (0.04)
Influence absolute diff.	-0.03* (0.02)	-0.06*	0.04 [-0.01; 0.09]	-0.05* (0.01)
Alter = Government actor	0.63* (0.25)	0.68	-0.46 [-1.09; 0.13]	1.04* (0.34)
Functional requirements				
Ego = Environmental NGO	0.88* (0.26)	0.99	-0.60 [-1.31; 0.09]	0.79* (0.17)
Same actor type	0.74* (0.22)	1.12*	1.17* [0.61; 1.70]	0.99* (0.23)
Endogenous dependencies				
Mutuality	1.22* (0.21)	1.00*		0.81* (0.25)
Outdegree popularity				0.95* (0.09)
This article is protected by copyright. All rights reserved.				
Twopaths				-0.04* (0.02)
GWIdegree (2.0)				3.42* (1.45)
GWESP (1.0)				0.58* (0.16)
GWOdegree (0.5)				8.42*

Substantive comparison

ERGMs may yield a better fit, but does the difference between models affect the substantive results? Table 1 shows the estimated coefficients for all four models. While many results are consistent across different models, there are several coefficients that differ in their magnitude, significance, and even direction across models. In other words, different modeling strategies sometimes result in different substantive conclusions, which renders model selection based on endogenous fit an important task.

For example, the **Business vs. NGO** effect can be confirmed in all of the network models while the logit model shows an insignificant effect due to unmodeled dependencies. This highlights our point that appropriate treatment of dependencies between observations can make a substantive difference.

On the other hand, there are profound differences between the three network models. For example, the **Alter = Government actor** model term has a positive but insignificant effect in the QAP model, a positively significant effect in the ERGM, and a negative and insignificant effect in the latent space model. It is unclear in principle which model is “right.” Yet, given that the dependencies have been most accurately captured by the ERGM, the positively significant effect in the ERGM seems most trustworthy. It is also in line with the theoretical prediction.

Another important observation is that the latent space model deviates in many cases from the results of the QAP and ERGM. This happens in cases where a model term is correlated with the latent position of a vertex. The latent space satisfies the triangle inequality: vertices with a path distance of g are placed closer to each other than vertices with a path distance of $g + 1$, for any g . The consequence is that central actors are placed centrally in the latent space because they have smaller average path distances to any other vertex. Examples are the Swiss Agency for the Environment, Forests and Landscape (BUWAL) and the Swiss Petrol Union (EV), which are at the center of the latent space (see the actual latent space in the last panel of Figure 4) and also have large indegree, closeness, and betweenness centrality scores. Since many government actors are central players, the Euclidean distances pick up the government covariate as part of the dependencies. The consequence of this trivariate correlation between distance in the latent space (or centrality in the network), the link receiver (alter) being a government actor, and link sender (ego) sending information to the receiver is that controlling for latent distances reverses the coefficient for the covariate (or makes it insignificant in

other cases). This is why we see a negative coefficient for the receiver being a government actor while the QAP and the ERGM show the opposite. Thus, while latent space models offer an easily interpretable visual representation of social distance (last panel of Figure 4), this may also cause interpretability issues when covariates are correlated with the latent space and the outcome network. Further, including random sender and receiver effects, as is possible in the standard model, does not rectify the problem. A potential remedy is to use specifications of the latent space that explicitly correct for the activity and prominence of vertices via a bilinear model (??), but that is not implemented in software as of this writing.

The Bayesian Information Criterion (BIC) at the bottom of the table indicates that the ERGM offers the most preferable combination of model fit and number of parameters. BIC can be compared across the different models because the same data are modeled. The QAP is slightly better than the logit, and the latent space model—even though it does a better job of capturing the dependencies than the logit and QAP—has the largest BIC value, because of its many parameters.

Overall, the models confirm most of the hypotheses. The analysis confirms earlier findings that preference (dis)similarity is rendered unimportant if transaction cost considerations in terms of joint forum participation (?) and structural properties of the network are properly taken into account. Moreover, we find that collaboration in policy networks is a function of conflicting actor roles, influence attribution, and functional requirements.

Concluding Thoughts on Method Selection

Given the advantages and disadvantages of the models considered here, which model is preferable? Table 2 summarizes the characteristics of the models. Four criteria, in no particular order, should be considered: parsimony, interpretability, theoretical focus, and ease of use.

With regard to parsimony, the latent space model is the least appealing because, in addition to the covariates, a separate parameter is estimated for each actor’s latent position times the number of dimensions of the latent space. This is reflected in the relatively bad BIC value reported in Table 1. The QAP is the best choice with respect to parsimony because only the twelve (exogenous covariate) parameters are used. With QAP, the data can be described using relatively little external information. The ERGM has additional parameters for the endogenous dependencies, but is still relatively

Criterion	QAP	ERGM	LSM
Operationalization of relational theories	–	+	–
Easy to specify and interpret	+	–	◦
Parsimonious (= few parameters needed)	+	◦	–
Reports standard errors	–	+	+
Avoids problems with oversensitivity of dependencies	+	+	–
Simulations do not suffer from omitted variable bias	–	+	+
Avoids numerical instability (= degeneracy)	+	–	+
Unbiased under arbitrary empirical distributions	–	+	+
Flexible with respect to outcome distribution	+	+	+
Temporal specifications are available	–	+	+
Spatial visualization and model-based clustering	–	–	+
Full-fledged statistical model	–	+	+
Availability in standard statistical software	+	+	+

Table 2: Comparison of Three Models for Inferential Network Analysis

parsimonious.

With regard to interpretability, the QAP fares worst because dependencies cannot be inspected at all. Moreover, one must be careful about bias when the network to be modeled is sparse and some predictors are collinear; such cases may be common because many empirical networks have low density. The latent space model allows one to visualize the latent space and interpret it substantively, and the ERGM provides parameters for each endogenous part of the data-generating process; both the latent space model and the ERGM allow simulation-based goodness-of-fit assessment. This is also possible with the QAP, but due to the way the QAP corrects for dependencies rather than modeling them, simulations based on a QAP model do not include any complex dependencies that may be important for the data-generating process. This is reflected by the relatively bad endogenous model fit in Figure 3. The case study also illustrates that interpretability of the latent space model can be limited in some situations where covariates have to be evaluated as predictors of edges independently of the latent space, at least when Euclidean distances are used or when a bilinear model without activity and popularity corrections is employed.

With respect to theoretical focus, the ERGM is the only real option if a substantive theory about the types of dependencies shall be tested; neither

the latent space model nor QAP allow for hypothesis testing of different, nuanced dependency structures. On the other hand, if the goal is to test substantive theory about the exogenous covariates only and the dependencies are merely a nuisance, QAP and the latent space model allow one to do this without having to worry about the structure of the dependencies. It would be presumptuous for us to suggest that one model is superior to the others. The choice must depend upon the aims of the research, the theory, and the nature of the data. To the extent that an analysis is mainly concerned with network interdependency as a threat to analysis or wants to get a “clean” test of the relationships between the outcome network and exogenous covariates, QAP is often sufficient. If an analyst wants a model that can account for omitted variable bias, be interpreted graphically, and can be used for more than simple dyadic connections, use of the latent space model would work well. Finally, the ERGM is potentially the most powerful model when network structure is of interest in addition to the covariates, but it is also the most difficult to apply from an end-user’s perspective.