



UNIVERSITY OF LEEDS

This is a repository copy of *Dependence matters: Statistical models to identify the drivers of tie formation in economic networks*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/203489/>

Version: Accepted Version

Article:

De Nicola, G., Fritz, C., Mehrl, M. orcid.org/0000-0002-5825-9256 et al. (1 more author) (Cover date: November 2023) Dependence matters: Statistical models to identify the drivers of tie formation in economic networks. Journal of Economic Behavior and Organization, 215. pp. 351-363. ISSN 0167-2681

<https://doi.org/10.1016/j.jebo.2023.09.021>

© 2023, Elsevier. This manuscript version is made available under the CC-BY-NC-ND 4.0 license <http://creativecommons.org/licenses/by-nc-nd/4.0/>. This is an author produced version of an article published in the Journal of Economic Behavior and Organization. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Dependence matters: Statistical models to identify the drivers of tie formation in economic networks

Giacomo De Nicola^{1,*}, Cornelius Fritz², Marius Mehrl³, Göran Kauermann¹

Department of Statistics, LMU Munich, Ludwigstr. 33, 80539 Munich, Germany¹

Department of Statistics, Pennsylvania State University, 100 Thomas Building,
16802 State College, PA, USA²

School of Politics and International Studies, University of Leeds, LS2 9JT Leeds,
United Kingdom³

Abstract

Networks are ubiquitous in economic research on organizations, trade, and many other areas. However, while economic theory extensively considers networks, no general framework for their empirical modeling has yet emerged. We thus introduce two different statistical models for this purpose – the Exponential Random Graph Model (ERGM) and the Additive and Multiplicative Effects network model (AME). Both model classes can account for network interdependencies between observations, but differ in how they do so. The ERGM allows one to explicitly specify and test the influence of particular network structures, making it a natural choice if one is substantively interested in estimating endogenous network effects. In contrast, AME captures these effects by introducing actor-specific latent variables affecting their propensity to form ties. This makes the latter a good choice if the researcher is interested in capturing the effect of exogenous covariates on tie formation without having a specific theory on the endogenous dependence structures at play. After introducing the two model classes, we showcase them through real-world applications to networks stemming from international arms trade and foreign exchange activity. We further provide full replication materials to facilitate the adoption of these methods in empirical economic research.

Keywords – Inferential Network Analysis, Network Data, Endogeneity, Arms Trade, Foreign Exchange Networks, Statistical Modeling

JEL classification – C20, C49, F14, F31, L14

Declaration of interest: none

1 Introduction

The study of networks has established itself as a central topic in economic research (Jackson, 2008). Within the broader context of the study of complex and interdependent systems (see e.g. Flaschel et al., 1997, 2007, 2018), networks can be defined as interconnected structures which can naturally be represented through graphs. In the economic literature, networks have been extensively considered from a theoretical perspective, with the primary goal of understanding how economic behavior is shaped by interaction patterns (Jackson and Rogers, 2007). Indeed, the adequate modelling of such interactions has been described as one of the main empirical challenges in economic network analysis (Jackson et al., 2017). Research in this direction on, e.g., organizations as networks, diffusion in networks, network experiments, or network games, is surveyed in Bramoullé et al. (2016), Jackson (2014), and Jackson et al. (2017). These theoretical advances find application in

*Corresponding author: giacomo.denicola@stat.uni-muenchen.de

many different fields in which network structures naturally arise, such as national and international trade, commercial agreements, firms’ organization, and collaboration activity. However, such advances have not yet been accompanied by a corresponding shift in the standard methods used to empirically validate them. Some recent contributions (see e.g. Atalay et al., 2011; Chaney, 2014; Morales et al., 2019) develop estimators tailored specifically to their network-based theoretical models, but more generally applicable modeling frameworks for the analysis of real-world network data have not yet emerged. Statistical methods specifically designed to empirically test theories where interdependencies arise from network structures, such as the Exponential Random Graph Model (ERGM), exist but are not yet widely used by economists. Jackson (2014), for instance, discusses ERGMs but argues that they “suffer from proven computational problems” (2014, p.76). Jackson et al. (2017) further explain that “it is practically impossible to estimate the likelihood of a given network at even a moderately large scale”, concluding that with ERGMs, “there is an important computational hurdle that must be overcome in working with data” (2017, p.85).

Contrasting this assessment, we argue that recent work in the realm of empirical network analysis provides robust and scalable methods with readily available implementations in the R statistical software (R Core Team, 2021). Computational issues thus do not represent an insurmountable barrier to employ robust inferential network methods anymore. In this paper, we demonstrate the effectiveness and usability of some of those methods by applying them to real economic data. We specifically focus on models which aim to capture the mechanisms leading to network formation, i.e. to measure how the probability of forming a tie is influenced by (a) nodal characteristics, (b) pairwise covariates, and (c) the rest of the network. In particular, our focus is on Exponential Random Graph Models (ERGM) (Robins et al., 2007) and Additive and Multiplicative Effect (AME) network models (Hoff, 2021), respectively implemented in the R packages **statnet** (Handcock et al., 2008) and **amen** (Hoff, 2015). We find these two model classes to be among the most promising ones for applications in the economic sciences, as they are well suited for answering two broad categories of research questions. The ERGM is an ideal fit if, based on economic theory, the researcher envisages a particular dependence structure for the existence of ties in the network at hand and wants to test whether their theory is corroborated by empirical data. On the other hand, AME, and more generally continuous latent variable models, are a good choice when the researcher is interested in capturing the effect of exogenous variables on tie formation without having prior knowledge on which endogenous network dependence mechanisms are at play. In this case, AME offers the possibility to estimate the effect of both nodal and pairwise covariates while simultaneously controlling for network effects, which may induce bias if ignored (see Lee and Ogburn, 2021). In addition, the estimated latent structure can provide insight on the underlying network mechanisms for which they are controlling.

The principal aim of this paper is to showcase ERGM and AME by focusing on their value for economic research. After introducing each model class, we demonstrate their empirical usage by respectively applying them to two relevant economic questions stemming from real-world networks. We first use the ERGM to model the international trade of major conventional weapons, where a directed tie exists if one country transfers arms to another. In line with Chaney (2014), network effects such as directed triadic closure (e.g. the positive impact of an increase in the volume of trade between countries A and B on the probability that country C, that already exports to A, starts exporting to B) are of explicit theoretical interest in this application, and the ERGM allows for

their proper specification and testing. We then make use of the AME model to study a historical network of global foreign exchange activity, where a directed edge is present if one country’s national currency is actively traded within the other country. AME allows us to estimate how relevant country features, such as per-capita gdp and the gold standard, and pairwise covariates, such as the distance between two countries and their reciprocal trade volume, influence tie formation, while controlling for network effects to provide unbiased estimates. We further compare the two model classes, weighing pros and cons of each approach and providing guidance on which tool is appropriate for applications to different empirical settings and research questions. Finally, in addition to a step-by-step analysis and interpretation of these application cases, we provide full replication code in our GitHub repository¹, allowing for seamless reproducibility. We, therefore, demonstrate the “off-the-shelf” applicability of these methods, and offer applied researchers a head-start in employing them to study substantive economic problems.

Our contribution is related to various strands of the growing literature on economic networks (e.g. Jackson and Rogers, 2007; Jackson, 2008; Bramoullé et al., 2016). Due to its focus on economic questions, our work differs from surveys in physics (Newman, 2003), statistics (Goldenberg et al., 2010), or political science (Cranmer et al., 2017). Several articles provide overviews and surveys of existing economic network models from a theoretical perspective (Jackson, 2014; Graham, 2015; Jackson et al., 2017; De Paula, 2020). None of these articles concentrates on discussing broadly applicable statistical modeling frameworks, such as ERGM and AME, from an empirical perspective. In this sense our paper is similar in spirit to van der Pol (2019) who, however, only focuses on ERGM, without comparing alternative approaches. Indeed, one of the goals of this paper is to shed light on the emerging AME model class (and, more generally, on latent variable network models) for future applications in the economic literature.

The remainder of the paper is structured as follows. Section 2 discusses existing literature and presents the mathematical and notational framework used to define and discuss networks throughout the paper. Section 3 introduces the ERGM and applies it to the international arms trade network. Section 4 is dedicated to AME and its application to the global foreign exchange network. Section 5 concludes the paper with a brief discussion on the two model classes, contrasting their different uses and highlighting pros and cons of each approach.

2 Economic Networks

2.1 Related literature

Even though network structures naturally arise in many aspects of economics and are subject of prominent research in the field, much of the previous literature has ignored the implied interdependencies, instead opting for regression models assuming ties to be independent conditional on the covariates (e.g. Anderson and Van Wincoop, 2003, Rose, 2004, Lewer and Van den Berg, 2008). This assumption is often unreasonable in practice. It would, for example, imply that Germany imposing economic sanctions on Russia is independent of Italy imposing sanctions on Russia, and, in the directed case, even of Russia imposing them on Germany itself. While no standard framework for the modeling of empirical network data has emerged in economics so far, a number of contributions in

¹<https://github.com/gdenicola/statistical-network-analysis-in-economics>

– or adjacent to – the field do make use of statistical network models. We shortly survey these works here to show that the models we present are indeed suitable for the analysis of economic data. Possibly the most obvious kind of economic network is the international trade network (see Chaney, 2014) and many of these studies accordingly seek to model the formation of trade ties. In this vein, two early studies (Ward and Hoff, 2007; Ward et al., 2013) apply latent position models to show that trade exhibits a latent network structure beyond what a standard gravity model can capture (see also Fagiolo, 2010; Dueñas and Fagiolo, 2013). More recently, numerous contributions have used the ERGM to explicitly theorize and understand network interdependence in the general trade (Herman, 2022; Liu et al., 2022; Smith and Sarabi, 2022) as well as the trade in arms (Thurner et al., 2019; Lebacher et al., 2021), patents (He et al., 2019), and services (Feng et al., 2021).

That being said, empirical research on economic networks is not limited to trade. Smith et al. (2019) use multilevel ERGMs to study a production network consisting of ownership ties between firms at the micro-level and trade ties between countries at the macro-level, while Mundt (2021) explores the European Union’s sector-level production network via ERGMs as well as an alternative methodology, the stochastic actor-oriented model (SAOM). The latter is another prominent tool in the realm of network analysis, which is suitable for modeling longitudinal network data. As we, in the interest of brevity, focus on models for static networks (i.e. networks that are observed only at one point in time), we do not treat the SAOM, and instead refer to Snijders (1996, 2017) for an introduction to the model class. Going back to empirical research on economic networks in the literature, Fritz et al. (2023) deploy ERGMs to investigate patent collaboration networks. Studies on foreign direct investments document network influences using latent position models (Cao and Ward, 2014), or seek to model them via extensions of the ERGM (Schoeneman et al., 2022). Finally, economists also study networks of interstate alliances and armed conflict (see e.g. Jackson and Nei, 2015; König et al., 2017), both of which have been modeled via ERGMs (Cranmer et al., 2012; Campbell et al., 2018) and AME (Dorff et al., 2020; Minhas et al., 2022). This short survey indicates that both ERGM and AME can be used to answer questions which are of substantive interest to economists.

2.2 Setup

Before introducing models for networks in which dependencies between ties are expected, we briefly introduce the mathematical framework for networks, as well as the necessary notation. Let $\mathbf{y} = (y_{ij})_{i,j=1,\dots,n}$ be the adjacency matrix representing the observed binary network, comprising n fixed and known agents (nodes). In this context, $y_{ij} = 1$ indicates an edge from agent i to agent j , while $y_{ij} = 0$ translates to no edge between the two. Since self-loops are not admitted for most studied networks, the diagonal of $\mathbf{y}$ is left unspecified or set to zero. Depending on the application, the direction of an edge can carry additional information. If it does, we call the network directed. In this article, we mainly focus on this type of networks. Also note that all matrix-valued objects are written in bold font for consistency. In addition to the network connections, we often observe covariate information on the agents, which can be at the level of single agents (e.g. the gdp of a country) or at the pairwise level (e.g. the distance between two countries). We denote covariates by $\mathbf{x}_1, \dots, \mathbf{x}_p$, and our goal is to specify a statistical model for $\mathbf{Y}$, that is the random variable corresponding to $\mathbf{y}$, conditional on $\mathbf{x}_1, \dots, \mathbf{x}_p$. A natural way to do this is to specify a probability distribution over the space of all possible networks, which we define by the set $\mathcal{Y}$. Two main characteristics differentiate our modeling endeavor from

classical regression techniques, such as Probit or logistic regression models. First, for most applications, we only observe one realization $\mathbf{y}$ from $\mathbf{Y}$, rendering the estimation of the parameters to characterize this distribution particularly challenging. Second, the entries of $\mathbf{Y}$ are generally co-dependent; thus, most conditional dependence assumptions inherent to common regression models are violated. Generally, we term mechanisms that induce direct dependence between edges to be endogenous, while all effects external to the modeled network, such as covariates, are called exogenous.

3 The Exponential Random Graph Model

The ERGM is one of the most popular models for analyzing network data. First introduced by Holland and Leinhardt (1981) as a model class that builds on the platform of exponential families, it was later extended with respect to fitting algorithms and more complex dependence structures (Lusher et al., 2012; Robins et al., 2007). We next introduce the model step-by-step to highlight its ability to progressively generalize by building on conditional dependence assumptions.

3.1 Accounting for dependence in networks

We begin with the simplest possible stochastic network model, the Erdős-Rényi-Gilbert model (Erdős and Rényi, 1959; Gilbert, 1959), where all edges are assumed to be independent and to have the same probability of being observed. In stochastic terms, each observed tie is then a realization of a binomial random variable with success probability π , which yields

$$\mathbb{P}_{\pi}(\mathbf{Y} = \mathbf{y}) = \prod_{i=1}^n \prod_{j \neq i} \pi^{y_{ij}} (1 - \pi)^{1-y_{ij}} \quad (1)$$

for the probability to observe $\mathbf{y}$. Evidently, model (1), which implies equal probability for all possible ties, is too restrictive to be applied to real world problems. In the next step, we, therefore, additionally incorporate covariates x_{ij} by letting π vary depending on those covariates, leading to edge-specific probabilities π_{ij} . Following the common practice in logistic regression, we parameterize the log-odds by $\log\left(\frac{\pi_{ij}}{1-\pi_{ij}}\right) = \boldsymbol{\theta}^{\top} x_{ij}$, where x_{ij} is a vector of exogenous statistics with the first entry set to 1 to incorporate an intercept, and get

$$\mathbb{P}_{\boldsymbol{\theta}}(\mathbf{Y} = \mathbf{y}) = \prod_{i=1}^n \prod_{j \neq i} \left(\frac{\exp\{\boldsymbol{\theta}^{\top} x_{ij}\}}{1 + \exp\{\boldsymbol{\theta}^{\top} x_{ij}\}} \right)^{y_{ij}} \left(\frac{1}{1 + \exp\{\boldsymbol{\theta}^{\top} x_{ij}\}} \right)^{1-y_{ij}}. \quad (2)$$

From (2), the analogy to standard logistic regression being a special case of generalized linear models (Nelder and Wedderburn, 1972) becomes apparent. The joint distribution of $\mathbf{Y}$ can be formulated in exponential family form, yielding

$$\mathbb{P}_{\boldsymbol{\theta}}(\mathbf{Y} = \mathbf{y} | x) = \frac{\exp\{\boldsymbol{\theta}^{\top} s(\mathbf{y})\}}{\kappa(\boldsymbol{\theta})}, \quad (3)$$

where $s(\mathbf{y}) = (s_1(\mathbf{y}), \dots, s_p(\mathbf{y}))$, $s_q(\mathbf{y}) = \sum_{i=1}^n \sum_{j \neq i} y_{ij} x_{ij,q} \forall q = 1, \dots, p$, with $x_{ij,q}$ as q -th entry in x_{ij} and $\kappa(\boldsymbol{\theta}) = \prod_{i=1}^n \prod_{j \neq i} (1 + \exp\{\boldsymbol{\theta}^{\top} x_{ij}\})$. In the jargon of exponential families, we term $s(\mathbf{y})$ sufficient statistics.




Figure 1: Illustration of directed edgewise-shared partner statistics for k agents. Circles represent agents, and black lines represent edges between them. The names follow **statnet** nomenclature: OTP = “Outgoing Two-Path”, ISP = “Incoming Shared Partner”, OSP = “Outgoing Shared Partner”, and ITP = “Incoming Two-Path”.

Newcomb (1979) observed that many observed networks exhibit complicated relational mechanisms, including reciprocity, which we can account for by extending the set of sufficient statistics. Under reciprocity, an edge Y_{ij} influences the probability of its reciprocal edge Y_{ji} to occur. Analyzing social networks, we would expect that the probability of agent i nominating agent j to be a friend is higher if agent j has nominated agent i as a friend. Holland and Leinhardt (1981) extended model (1) to such settings with the so-called p_1 model. To represent reciprocity, we assume dyads, each of them defined by (Y_{ij}, Y_{ji}) , to be independent of one another, which again yields an exponential family distribution similar to (3) with sufficient statistics that count the number of mutual ties ($s_{\text{Mut}}(\mathbf{y}) = \sum_{i < j} y_{ij} y_{ji}$), of edges ($s_{\text{Edges}}(\mathbf{y}) = \sum_{i=1}^n \sum_{j \neq i} y_{ij}$), and the in- and out-degree statistics for all degrees observed in the networks². Agents’ in- and out-degrees are their number of incoming and outgoing edges, and relate to their relative position in the network (Wasserman and Faust, 1994).

Next to reciprocity, another important endogenous network mechanism is transitivity, originating in the structural balance theory of Heider (1946) and adapted to binary networks by Davis (1970). Transitivity affects the clustering in the network, implying that a two-path between agents i and j , i.e. $y_{ih} = y_{hj} = 1$ for some other agent h , affects the edge probability of Y_{ij} . Put differently, Y_{ij} and Y_{kh} are assumed to be independent iff $i, j \neq k$ and $i, j \neq h$. Frank and Strauss (1986) proposed the Markov model to capture such dependencies. For this model, the sufficient statistics are star-statistics, which are counts of sub-structures in the network where one agent has (incoming and outgoing) edges to between 0 and $n - 1$ other agents, and counts of triangular structures. If the network is directed it is possible to define different types of triangular structures, as depicted in Figure 1.

3.2 Extension to general dependencies

Starting from the Erdős-Rényi-Gilbert model, which is a special case of a generalized linear model, we have consecutively allowed for more complicated dependencies between edges, resulting in the Markov graphs of Frank and Strauss (1986). Over this course, we showed that each model can be stated in exponential family form, characterized by a particular set of sufficient statistics. We now make this more explicit to allow for more general dependence structures, and specify a probabilistic model for $\mathbf{Y}$ directly through

²We provide more details on this derivation in the Supplementary Material.

the sufficient statistics³. Wasserman and Pattison (1996) introduced this model as

$$\mathbb{P}_{\boldsymbol{\theta}}(\mathbf{Y} = \mathbf{y}) = \frac{\exp\{\boldsymbol{\theta}^\top s(\mathbf{y})\}}{\kappa(\boldsymbol{\theta})}, \quad (4)$$

where $\boldsymbol{\theta}$ is a p -dimensional vector of parameters to be estimated, $s(\mathbf{y})$ is a function calculating the vector of p sufficient statistics for network $\mathbf{y}$, and $\kappa(\boldsymbol{\theta}) = \sum_{\tilde{\mathbf{y}} \in \mathcal{Y}} \exp\{\boldsymbol{\theta}^\top s(\tilde{\mathbf{y}})\}$ is a normalizing constant to ensure that (4) sums up to one over all $\mathbf{y} \in \mathcal{Y}$. To estimate $\boldsymbol{\theta}$, Handcock (2003) adapted the Monte Carlo Maximum Likelihood technique of Geyer and Thompson (1992), approximating the logarithmic likelihood ratio of $\boldsymbol{\theta}$ and a fixed $\boldsymbol{\theta}_0$ via Monte Carlo quadrature (see Hunter et al., 2012, for an in-depth discussion).

A problem often encountered when fitting model (4) to networks is degeneracy (Handcock, 2003; Schweinberger, 2011). Degenerate models are characterized by probability distributions that put most probability mass either on the empty or on the full network, i.e., where either all or no ties are observed. To detect this behavior, one can use a goodness-of-fit procedure where observed network statistics are compared to statistics of networks simulated under the estimated model (Hunter et al., 2008). To address it, Snijders et al. (2006) and Hunter and Handcock (2006) propose weighted statistics that, in many cases, have better empirical behavior. Degeneracy commonly affects model specifications encompassing statistics for triad counts and multiple degree statistics. For in-degree statistics, we would thus incorporate the geometrically-weighted in-degree,

$$GWIDEG(\mathbf{y}, \alpha) = \exp\{\alpha\} \sum_{k=1}^{n-1} \left(1 - (1 - \exp\{-\alpha\})^k\right) IDEG_k(\mathbf{y}), \quad (5)$$

where $IDEG_k(\mathbf{y})$ is the number of agents in the studied network with in-degree k and α is a fixed decay parameter. One can substitute $IDEG_k(\mathbf{y})$ in (5) with the number of agents with a specific out-degree, $ODEG_k(\mathbf{y})$, to capture the out-degree distribution. We term these statistics geometrically weighted since the weights in (5) are a geometric series⁴. A positive estimate implies that an edge from a low-degree agent is more likely than an edge from a high-degree agent, resulting in a decentralized network. If, on the other hand, the corresponding coefficient is negative, one may interpret it as an indicator for a centralized network.

To capture clustering, we have to define the distribution of edgewise-shared partners (ESP). This distribution is defined as the relative frequency of edges in the network with a specific number of k shared partners, that we denote by $ESP(\mathbf{y})$ for $k \in \{1, \dots, n-2\}$. As shown in Figure 1, various versions of edgewise-shared partner statistics can be found in directed networks, depending on the direction of the edges between the three agents involved. Geometrically weighted statistics can be stated for them in a similar manner as for degree statistics. For example, for the outgoing two-path (OTP, see Figure 1a), this is

$$GWOTP(\mathbf{y}, \alpha) = \exp\{\alpha\} \sum_{k=1}^{n-2} \left(1 - (1 - \exp\{-\alpha\})^k\right) OTP_k(\mathbf{y}). \quad (6)$$

³Alternatively, (4) can also be derived as the equilibrium distribution of a strategic game where players myopically reassess and update their links to optimize their utility in the network (see Mele, 2017; Boucher and Mourifié, 2017).

⁴Geometrically weighted statistics require setting the decay parameter α . We set $\alpha = \log(2)$, though it can also be estimated as an additional parameter given sufficient data (Hunter and Handcock, 2006).

In this case, a positive coefficient indicates that sharing ties with third actors increases the probability of observing an event between two agents.

Along with capturing endogenous network statistics, it is also possible to extend the ERGM framework to include the temporal dimension, that is, to model longitudinal network data. This is done quite naturally through use of a Markov assumption on the temporal dependence of subsequently observed networks, giving rise to the Temporal Exponential Random Graph Model (TERGM). As we here focus on static networks, we do not cover this in depth, and refer to Hanneke et al. (2010) for an introduction to the TERGM, and to Fritz et al. (2020) for a more general discussion on temporal extensions to the model class.

In summary, the ERGM allows to account for network dependencies via explicitly specifying them in $s(\mathbf{y})$. A large variety of potential network statistics, such as those given in (5) and (6), can be included in $s(\mathbf{y})$, enabling to test for their influence in the formation of the observed network. By allowing for this explicit inclusion and testing of network statistics, the ERGM requires researchers to at least have an implicit theory regarding what types of network dependence should exist in the network they study. Without such theory to guide the selection of network statistics, the range of potential network dependencies, and corresponding statistics, is virtually endless⁵. As a result, the ERGM is best suited for research questions that explicitly concern interdependencies within the network. If these interdependencies are, instead, only a potential source of bias the researcher wants to control for, the AME model (introduced in Section 4) may be a better fit.

3.3 Application to the international arms trade network

We next make use of the ERGM to analyze the international arms transfer network. Recent studies on trade in Major Conventional Weapons (MCW), such as fighter aircraft or tanks, not only emphasize its networked nature, but also argue that this very nature is of substantive theoretical interest (Thurner et al., 2019; Fritz et al., 2021). In line with Chaney (2014), triadic trade structures are held to reveal information regarding the participants' economic and security interests. Explicitly modeling these structures allows us to test hypotheses regarding their effects on further arms transfers. Accordingly, we seek to model the network of international arms transfers in the year 2018, where countries are nodes and a directed edge indicates MCW being delivered from country i to country j . Our interest here mainly lies in uncovering the network's endogenous mechanisms. MCW trade data come from SIPRI (2021), and the resulting network is depicted in Figure 2, obtained using the Yifan Hu force-directed graph drawing algorithm (Hu, 2005) with the software Gephi (Bastian et al., 2009).

For estimating the parameters characterizing the ERGM, we use the R package `ergm` (Handcock et al., 2008). Since evaluating $\kappa(\boldsymbol{\theta})$ from (4) necessitates calculating the sum of $|\mathcal{Y}| = 2^{n(n-1)}$ terms, we rely on MCMC approximations thereof to obtain the maximum likelihood estimates (see Handcock, 2003 and Hummel et al., 2012 for additional information on this topic). As discussed above, the ERGM allows us to use both exogenous (node-specific and pair-specific) attributes as well as endogenous structures to model the network of interest. Here, we select both types of covariates based on existing studies on the arms trade (Thurner et al., 2019; Fritz et al., 2021). In addition to an edges term, which corresponds to the intercept in standard regression models, we include importers'

⁵For a survey of possible endogenous terms, see Morris et al. (2008).




Figure 2: Illustration of the international arms trade network in 2018. Countries are labeled by their ISO 3166-1 codes, and a directed edge from node i to node j indicates major conventional weapons being delivered from country i to country j .

and exporters' logged GDP, whether they share a defense pact, their absolute difference in "polity" scores (a type of democracy index), and their geographical distance⁶. We lag these covariates by three years, reflecting the median time between order and delivery for MCW delivered in 2018⁷. More importantly, for the purpose of demonstrating how to model network data with the ERGM, we specify five endogenous network terms. In- and out-degree (IDEG and ODEG) measure, respectively, importers' and exporters' trade activity, and thus capture whether highly active importers and exporters are particularly attractive trading partners, or if they are instead less likely to form additional trade ties. Moreover, we specify a reciprocity term to capture whether countries tend to trade MCW uni- or bidirectionally. We further include two types of triadic structures, which represent transitivity and a shared supplier between countries i and j . The transitivity term counts how often country i exports arms to j while i exports to k , which in turn exports to j , thus capturing i 's tendency to directly trade with j if they engage in indirect trade (OTP, see Figure 1a). In contrast, the shared supplier term counts how often country i sends arms to j while both import weapons from a shared supplier k (ISP, see Figure 1b). Note that, given the issue of degeneracy discussed above, we use geometrically weighted versions of all endogenous statistics except reciprocity. Finally, we include a repetition term capturing whether arms transfer dyads observed in 2018 had already occurred in any of the three previous years. Results of this ERGM, as well as, for comparison's sake, a logistic regression that includes the same exogenous covariates but does not capture

⁶Data for these covariates come from the **peacesciencer** package (Miller, 2022).

⁷We use the median as the distribution of times between order and delivery is quite skewed. As shown in the Supplementary Materials, our substantive results remain unchanged when using 4- and 5-year lags instead, which reflect the average time between order and delivery. In particular, the ERGM outperforms the logistic regression model regardless of lag choice.

	ERGM	Logit
Intercept	−15.356 (2.017)***	−28.197 (1.731)***
Repetition	3.254 (0.141)***	3.957 (0.141)***
Distance	−0.081 (0.087)	−0.239 (0.088)**
Abs. Diff. Polity	−0.001 (0.010)	−0.003 (0.012)
Alliance	0.350 (0.207)	0.209 (0.207)
log-GDP (Sender)	0.300 (0.050)***	0.588 (0.045)***
log-GDP (Receiver)	0.166 (0.049)***	0.355 (0.039)***
Mutual	−0.311 (0.438)	
GWIDEG	−1.478 (0.296)***	
GWODEG	−2.848 (0.296)***	
GWOTP	−0.146 (0.104)	
GWISP	0.210 (0.083)*	
AIC	1769.718	1891.984
BIC	1866.053	1948.179
Log Likelihood	−872.859	−938.992

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 1: Estimated coefficients and standard errors (in parentheses) of the ERGM and the logistic regression model for the international arms trade network in 2018.

any of the endogenous network structures, are presented in Table 1. These results can be compared directly, as, just like in a logistic regression model, coefficients in the ERGM indicate the additive change in the log odds of a tie occurring in association with a unit change in the respective variable. In this sense, the logistic regression model can be viewed as a special case of the ERGM in which the network effects are omitted. From the table, we can see how the two models differ both in their in-sample fit, as captured by AIC and BIC⁸, as well as in the substantive effects they identify for the exogenous covariates. The repetition coefficient is positive and statistically significant in both models, but differs substantially in its size. An arms transfer edge having occurred at least once in 2015-17 increases the log odds of it occurring also in 2018 by 3.96 in the Logit, but only by 3.26 in the ERGM. Similarly, both models agree that the log odds of an arms transfer occurring increase with the economic size of the sender and receiver, as captured by their respective GDPs, but the coefficients retrieved by the Logit are approximately double the size of those in the ERGM, thus attributing more explanatory power to them. Also in this vein, the effect of the geographical distance between sender and receiver is three times as large in the logistic regression as in the ERGM and, while statistically significant in the former, indistinguishable from zero in the latter. Finally, both models report small and statistically insignificant effects for countries' polity difference and alliance ties. Taken together, however, there are clear, substantively meaningful differences in the effect sizes

⁸As shown in the Supplementary Material, the ERGM also outperforms the Logit model when assessing their respective areas under the receiver-operator and precision-recall curves. In line with Hunter and Handcock (2006), one could also calculate the likelihood ratio test statistic from the log likelihoods reported in Table 1 for the same purpose.

and, in the case of geographical distance, even statistical significance of the coefficients that the ERGM and Logit recover for the exogenous covariates.

Furthermore, three of the endogenous statistics included in the ERGM exhibit statistically significant effects on the probability of arms being traded. The results for in- and out-degree replicate the finding by Thurner et al. (2019), showing that highly active importers and exporters are less likely to form additional trade ties. In the ERGM, coefficients can also be interpreted at the global level, in addition to the edge-level interpretation given above. The shared supplier term having a (statistically significant) positive coefficient indicates, at the edge level, that an exporter is more likely to transfer weapons to a potential receiver if both of them import arms from the same source. Globally, on the other hand, the same coefficient means that the observed network exhibits more shared supplier configurations – where country i sends weapon to j while both receive arms from k – than would be expected in a random network of the same size. On the whole, the results presented in Table 1 offer an example for the striking differences that modeling network structures (instead of assuming them away) can make. The ERGM and Logit, while identical in their non-network covariates, report substantively different effects for these covariates, and, in the ERGM, network effects are also found to drive the formation of arms transfer edges.

4 The Additive and Multiplicative Effect Network Model

4.1 Latent variable network models

Another way to account for network dependencies is by making use of latent variables. Models within this class assume that latent variables $\mathbf{Z}_i$ are associated with each node i . Depending on the type of model, these latent variables can either be discrete (e.g. indicating group memberships for each node) or continuous, and affect the connection probability in different ways (Matias and Robin, 2014). An early (but still popular) approach in this direction is the stochastic blockmodel, which assumes that each agent possesses a latent, categorical class (or group membership). Nodes within each class are assumed to be stochastically equivalent in their connectivity behavior, meaning that the probability of two nodes to connect depends solely on their group memberships (Holland et al., 1983; De Nicola et al., 2022). This family of models is attractive due to its simplicity in detecting and describing subgroups of nodes in networks. In many applications, however, discrete groupings fail to adequately represent the observed data, as agents behave more heterogeneously. Moving from discrete to continuous latent variable network models, another prominent approach is the latent distance model. The latter postulates that agents are positioned in a latent Euclidean “social space”, and that the closer they are within it, the more likely they are to form ties (Hoff et al., 2002). More precisely, the classical latent distance model specifies the probability of observing an edge between nodes i and j , conditional on $\mathbf{Z}$, through

$$\mathbb{P}_{\theta}(Y_{ij} = 1|\mathbf{Z}) = \frac{\exp\{\boldsymbol{\theta}^{\top} \mathbf{x}_{ij} - \|\mathbf{z}_i - \mathbf{z}_j\|\}}{1 + \exp\{\boldsymbol{\theta}^{\top} \mathbf{x}_{ij} - \|\mathbf{z}_i - \mathbf{z}_j\|\}}, \quad (7)$$

where $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_n)$ denotes the latent positions of the nodes in the d -dimensional latent space, and $\boldsymbol{\theta}$ is the coefficient vector for the covariates $\mathbf{x}_{ij}$. The latent positions

Z are assumed to originate independently from a spherical Gaussian distribution, i.e. $Z \sim N_d(0, \tau^2 \mathbf{I}_d)$, where $\mathbf{I}_d$ indicates a d -dimensional identity matrix.

Latent distance models are particularly attractive for social networks in which triadic closure plays a major role, and where nodes with similar characteristics tend to form connections with each other (i.e. homophilic networks, see Rivera et al., 2010). It is also possible to add nodal random effects to the model, to control for agent-specific heterogeneity in the propensity to form edges (Krivitsky et al., 2009). The model then becomes

$$\mathbb{P}_\theta(Y_{ij} = 1|Z, a, b) = \frac{\exp\{\theta^\top x_{ij} - \|z_i - z_j\| + a_i + b_j\}}{1 + \exp\{\theta^\top x_{ij} - \|z_i - z_j\| + a_i + b_j\}}, \quad (8)$$

where $a = (a_1, \dots, a_n)$ and $b = (b_1, \dots, b_n)$ are node-specific sender and receiver effects that account for the individual agents' propensity to form ties, with $a \sim N_n(0, \tau_a^2 \mathbf{I}_n)$ and $b \sim N_n(0, \tau_b^2 \mathbf{I}_n)$.

Despite its advantages and its fairly simple interpretation, a Euclidean latent space is unable to effectively approximate the behavior of networks where nodes that are similar in terms of connectivity behavior are not necessarily more likely to form ties (Hoff, 2008), such as, e.g., many networks of amorous relationships (Ghani et al., 1997; Bearman et al., 2004). More generally, the latent distance model tends to perform poorly for networks in which stochastic equivalence does not imply homophily and triadic closure, i.e., when nodes which behave similarly in terms of connectivity patterns towards the rest of the network do not necessarily have a higher probability of being connected among themselves. This is often the case in economics, where real-world networks can exhibit varying degrees and combinations of stochastic equivalence, triadic closure and homophily. Moreover, it is often *a priori* unclear which of these mechanisms are at play in a given observed network. In this context, agent-specific multiplicative random effects instead of the additive latent positions allow for simultaneously representing all these patterns (Hoff, 2005). Further developments of this innovation have led to the modern specification of the Additive and Multiplicative Effects network model (AME, Hoff, 2011), which, from a matrix representation perspective, generalizes both the stochastic blockmodel and the latent distance model (Hoff, 2021).

4.2 AME: Motivation and framework

The AME approach can be motivated by considering that network data often exhibit first-, second-, and third-order dependencies. *First-order effects* capture agent-specific heterogeneity in sending (or receiving) ties within a network. For example, in the case of companies and legal disputes, first-order effects can be viewed as the propensity of each firm to initiate (or be hit by) legal disputes. *Second-order effects*, i.e., reciprocity, describe the statistical dependency of the directed relationship between two agents in the network. In the previous example, this effect can be described as the correlation between (a) company i initiating a legal dispute against company j and (b) j doing the same towards i . Of course, second-order effects can only occur in directed networks. *Third-order effects* are described as the dependency within triads, defined as the connections between three agents, and relate to the triangular statistics previously illustrated in Figure 1. How likely is it that “a friend of a friend is also my friend”? Or, returning to the previous example: given that i has legal disputes with j and k , how likely are disputes to occur between j and k ?

The AME network model is designed to simultaneously capture these three orders

of dependencies. More specifically, it extends the classical (generalized) linear modeling framework by incorporating extra terms into the systematic component to account for them. In the case of binary network data, we can make use of the Probit AME model. As is well known, the classical Probit regression model can be motivated through a latent variable representation in which y_{ij} is the binary indicator that some latent normal random variable, say $L_{ij} \sim \mathcal{N}(\boldsymbol{\theta}^\top x_{ij}, \sigma^2)$, is greater than zero (Albert and Chib, 1993). But an ordinary Probit regression model assumes that L_{ij} , and thus the binary indicators (edges) y_{ij} , are independent, which is generally inappropriate for network data. In contrast, the AME Probit model specifies the probability of a tie y_{ij} from agent i to agent j , conditional on a set of latent variables W , as

$$\mathbb{P}(Y_{ij} = 1|W) = \Phi(\boldsymbol{\theta}^\top x_{ij} + e_{ij}), \quad (9)$$

where Φ is the cumulative distribution function of the standard normal distribution, $\boldsymbol{\theta}^\top x_{ij}$ accommodates the inclusion of dyadic, sender, and receiver covariates, and e_{ij} can be viewed as a structured residual, containing the latent terms in W to account for the network dependencies described above. In the directed case, e_{ij} is composed as

$$e_{ij} = a_i + b_j + u_i v_j + \varepsilon_{ij}. \quad (10)$$

In this context, a_i and b_j are zero-mean additive effects for sender i and receiver j accounting for first-order dependencies, jointly specified as

$$(a_1, b_1), \dots, (a_n, b_n) \stackrel{\text{i.i.d.}}{\sim} N_2(0, \Sigma_1), \quad \text{with} \quad \Sigma_1 = \begin{pmatrix} \sigma_a & \sigma_{ab} \\ \sigma_{ab} & \sigma_b \end{pmatrix}. \quad (11)$$

The parameters σ_a and σ_b measure the variance of the additive sender and receiver effects, respectively, while σ_{ab} relates to the covariance between sender and receiver effects for the same node. Going back to (10), ε_{ij} is a zero-mean residual term which accounts for second order dependencies, i.e. reciprocity. More specifically, it holds that

$$\{(\varepsilon_{ij}, \varepsilon_{ji}) : i < j\} \stackrel{\text{i.i.d.}}{\sim} N_2(0, \Sigma_2), \quad \text{with} \quad \Sigma_2 = \sigma^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}, \quad (12)$$

where σ^2 denotes the error variance and ρ determines the correlation between ε_{ij} and ε_{ji} , thus quantifying the tendency towards reciprocity. Finally, u_i and v_j in (10) are d -dimensional multiplicative sender and receiver effect vectors that account for third-order dependencies, and for which $(u_1, v_1), \dots, (u_n, v_n) \sim \mathcal{N}_{2d}(0, \Sigma_3)$ holds.

As noted above, AME is able to represent a wide variety of network structures, generalizing several other latent variable model classes. This generality comes at the price of a high level of complexity for the estimated latent structure. This can make the model class a sub-optimal choice if one wants to interpret the latent structure with respect to, e.g., clustering. On the other hand, its flexibility makes it an ideal fit when the underlying network dependencies are unknown, and the researchers' interest mainly lies in evaluating and interpreting the effect of dyadic and nodal covariates on tie formation while controlling for network effects. This strength has led to AME being used for several applications of this type (Koster, 2018; Minhas et al., 2019, 2022; Dorff et al., 2020). We next showcase the AME framework by applying it to the world foreign exchange activity network as of 1900, originally introduced and studied by Flandreau and Jobst (2005, 2009). This application highlights how using AME instead of classical regression can allow us to reconsider existing, influential answers to relevant questions via replication.




Figure 3: Illustration of the global foreign exchange activity network in 1900. Countries are labeled by their ISO 3166-1 codes, and an edge from node i to node j indicates active trading of the currency from country j within a financial center of country i .

4.3 Application to the global foreign exchange activity network

In 1900, every financial center featured a foreign exchange market where bankers bought and sold foreign currency against the domestic one. Foreign exchange market activity was monitored in local bulletins, which allowed Flandreau and Jobst (2005) to collect a global dataset with all currencies used in the world at that time. In the resulting network structure, laid out in Figure 3, countries are nodes, and a (directed) edge from country i to country j occurs if the currency of country j was actively traded in at least one financial center within country i . From the graph representation, laid out using a variant of the Yifan Hu force-directed graph drawing algorithm (Hu, 2005), we observe that the most actively traded currencies at the time belonged to large European economies, such as Great Britain, France and Germany. To determine the drivers of currency adoption, Flandreau and Jobst (2009) model this network as a function of several covariates by employing ordinary binary regression. As we show, it is possible to use AME to pursue the same goal while taking network dependencies into account.

We specify the AME model as in (9), using directed edges y_{ij} as response variable. The nodal covariates we use, sourced from and described in detail in the replication materials of Flandreau and Jobst (2009), are (log)per-capita GDP, democracy index score, coverage of foreign currencies traded in the country, and an indicator of whether the country's currency was on the gold standard. We also include, as dyadic covariates, the distance between two countries as well as their total trade volume. As specified in (10), the structured residual term e_{ij} comprises additive effects a_i and b_j for each node, which capture country-specific propensities to send and receive ties, respectively. Multiplicative effects u_i and v_j are included to account for third order dependencies. We here set the dimensionality of the multiplicative effects to two, which we assume to be sufficient given the relatively small size of the network.

To estimate the AME model, we make use of the R package `amen` (Hoff, 2015). As the likelihood involves intractable integrals arising from the combination of the transformation

and dependencies induced by the the model, closed form solutions are not available. The package thus uses reasonably standard Gibbs sampling algorithms to provide Bayesian inference on the model parameters. More details on the estimation routine can be found in Hoff (2021).

The results of the analysis, as well as, for comparison’s sake, a Probit regression including the same covariates but ignoring network dependencies, are displayed in Table 2. Note that the classical Probit regression model can be seen as a special case of AME Probit in which both additive and multiplicative node-specific effects are omitted. Additional model diagnostics and goodness of fit measures, together with the estimated variance and covariance parameters, are provided in the Supplementary Material. The estimated coefficients (for both models) can be interpreted as in standard Probit regression: For the nodal covariate per-capita GDP, for example, a unit increase in the log-per-capita GDP for country i corresponds to a decrease of 0.453 in the linear predictor, therefore negatively influencing the expected probability of the country to *send* a tie. The same unit increase in the log-per-capita-gdp for country i corresponds to an increase of 0.426 in the linear predictor, and has therefore a positive impact on the expected probability of that country to *receive* a tie. In the case of a dyadic covariate, such as distance, a unit increase in distance between two countries leads to a decrease of 1.019 in the linear predictor, resulting in a decrease in the expected probability of the two countries to form a tie in either direction. Overall, we find that the principal drivers of the formation of a tie between i and j are the magnitude of the foreign exchange coverage of the two countries involved, the distance between them, and their reciprocal trade volume. These results correspond to the thesis of Kindleberger (1967) and to Flandreau and Jobst (2009), who suggest that the most important determinants of international adoption for a currency are size and convenience of use. At the same time, we note that, as for the ERGM in the arms trade example, the results of the Probit and AME model differ in several regards. In particular, several effects are statistically significant in the Probit but not significant in the AME model. Indeed, unacknowledged network dependence can cause downward bias in the estimation of standard errors, leading to spurious associations (Lee and Ogburn, 2021). This finding once again highlights how accounting for network dependencies can make a difference when it comes to the substantive results.

As a final note, we add that in this case we went with AME over ERGM as our interest lies in answering the research questions addressed by Flandreau and Jobst (2005, 2009), that is assessing the effect of the exogenous covariates in Table 2 on tie formation. AME allows us to do that without specifying the configuration of the endogenous network mechanisms at play, which are instead accounted for through the imposed latent structure. If, on the other hand, the researcher expects some specific network effects to play a role, and wishes to test for their presence and measure their influence on network formation, the ERGM may be a better tool. The latter model class can, for example, directly answer questions such as “Does the fact that both countries A and B trade the currency of country C influence the probability of A and B to be connected? And if so, to what extent?”. AME, on the other hand, is limited to accounting for those effects via the latent variables, without explicitly identifying them, to provide unbiased inference for the covariate effects. The choice between the two model classes is thus a matter of what assumptions can be made about the network and where the researcher’s interest lies.

		AME	Classical Probit
Sender	Intercept	−4.845 (5.310)	−3.211 (1.580)*
	Gold standard	−0.629 (0.397)	−0.354 (0.155)*
	log-GDP per-capita	−0.453 (0.419)	−0.259 (0.152)
	Democracy index	−0.033 (0.064)	−0.025 (0.026)
	Currency coverage	1.418 (0.405)***	0.470 (0.137)***
Receiver	Gold standard	−0.599 (0.667)	−0.468 (0.191)*
	log-GDP per-capita	0.426 (0.703)	0.240 (0.159)
	Democracy index	0.121 (0.102)	0.066 (0.019)***
	Currency coverage	2.734 (0.691)***	1.363 (0.181)***
Dyadic	Distance	−1.019 (0.151)***	−0.471 (0.064)***
	log-trade volume	0.488 (0.081)***	0.346 (0.036)***

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 2: Estimated coefficients and related standard errors (in parentheses) for the AME model and the corresponding Probit model for the global foreign exchange activity network in 1900.

5 Conclusion

Complex dependencies are ubiquitous in the economic sciences (Chiarella et al., 2005; Flaschel et al., 2008), and many economic interactions can be naturally perceived as networks. This area of research has thus received considerable interest in recent years. However, this attention has not yet been accompanied by a corresponding general take-up of empirical research methods tailored towards networks. Instead, researchers either develop their own estimators to reproduce the features of their theoretical network models, or use standard regression methods that assume conditional independence of the edges in the network. Against this background, this paper seeks to provide a hands-on introduction to two statistical models which account for network dependencies, namely the Exponential Random Graph Model (ERGM) and the Additive and Multiplicative Effects network model (AME). These two classes serve different purposes: While the ERGM is most appropriate when explicitly interested in testing the effects of endogenous network structure, the AME model allows one to control for network dependencies while substantively focusing on estimating the effects of exogenous covariates of interest. We present the statistical foundations of both models, and demonstrate their applicability to economic networks through examples in the international arms trade and foreign currency exchange, showing that modeling network dependencies can alter the substantive results of the analysis. We, moreover, provide the full data and code necessary to replicate these exemplary applications. We explicitly encourage readers to use these replication materials to get started with analyzing economic networks via ERGM and AME, beginning with the examples covered here to then transfer the code and methods to their own research.

We especially want to encourage the use of such methods as not accounting for interdependence between observations when it exists can lead to biased estimates and spurious findings. Our two applications demonstrate that this bias can result in very different em-

pirical results, and thus affect substantive conclusions. It is therefore vital to account for network structure when studying interactions between economic agents such as individuals, firms, or countries, regardless of whether one is substantively interested in this structure. As shown by Lee and Ogburn (2021), our applications are just two examples of how unaccounted dependence in the observed data may lead to spurious findings.

At the same time, this paper can only serve as an introduction to statistical network data analysis in economics. We covered two general frameworks in this realm, but, in the interest of brevity, focused only on their simplest versions that apply to networks observed at only one time point and with binary edges. However, both frameworks have been extended to cover more general settings. For the ERGM, there are extensions for longitudinal data (Hanneke et al., 2010), distinguishing between edge formation and continuation (Krivitsky and Handcock, 2014), as well as to settings where edges are not binary but instead count-valued or signed (Krivitsky, 2012; Fritz et al., 2022). As for AME, approaches for longitudinal networks are described by Minhas et al. (2016), while versions for undirected networks as well as for non-binary network data are presented by Hoff (2021). Both the ERGM and the AME frameworks are thus flexible enough to cover a wide array of potential economic interactions. We believe that increasingly adopting these methods will, in turn, aid our understanding of these interactions.

References

- Albert, J. H. and S. Chib (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association* 88(422), 669–679.
- Anderson, J. E. and E. Van Wincoop (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review* 93(1), 170–192.
- Atalay, E., A. Hortacsu, J. Roberts, and C. Syverson (2011). Network structure of production. *Proceedings of the National Academy of Sciences* 108(13), 5199–5202.
- Bastian, M., S. Heymann, and M. Jacomy (2009). Gephi: An open source software for exploring and manipulating networks. *Third International AAAI Conference on Weblogs and Social Media*, 361–362.
- Bearman, P. S., J. Moody, and K. Stovel (2004). Chains of affection: The structure of adolescent romantic and sexual networks. *American journal of sociology* 110(1), 44–91.
- Boucher, V. and I. Mourifié (2017). My friend far, far away: A random field approach to exponential random graph models. *Econometrics Journal* 20(3), S14–S46.
- Bramoullé, Y., A. Galeotti, and B. W. Rogers (2016). *The Oxford Handbook of the Economics of Networks*. Oxford University Press.
- Campbell, B., S. Cranmer, and B. Desmarais (2018). Triangulating war: network structure and the democratic peace. *arXiv preprint*. arXiv:1809.04141.
- Cao, X. and M. D. Ward (2014). Do democracies attract portfolio investment? Transnational portfolio investments modeled as dynamic network. *International Interactions* 40(2), 216–245.
- Chaney, T. (2014). The network structure of international trade. *American Economic Review* 104(11), 3600–3634.
- Chiarella, C., P. Flaschel, and R. Franke (2005). *Foundations for a disequilibrium theory of the business cycle: Qualitative analysis and quantitative assessment*. Cambridge University Press.
- Cranmer, S. J., B. A. Desmarais, and J. H. Kirkland (2012). Toward a network theory of alliance formation. *International Interactions* 38(3), 295–324.
- Cranmer, S. J., P. Leifeld, S. D. McClurg, and M. Rolfe (2017). Navigating the range of statistical tools for inferential network analysis. *American Journal of Political Science* 61(1), 237–251.
- Davis, J. A. (1970). Clustering and hierarchy in interpersonal relations: Testing two graph theoretical models on 742 sociomatrices. *American Sociological Review* 35(5), 843–851.
- De Nicola, G., B. Sischka, and G. Kauermann (2022). Mixture models and networks: The stochastic blockmodel. *Statistical Modelling* 22(1-2), 67–94.
- De Paula, Á. (2020). Econometric models of network formation. *Annual Review of Economics* 12, 775–799.

- Dorff, C., M. Gallop, and S. Minhas (2020). Networks of violence: Predicting conflict in Nigeria. *Journal of Politics* 82(2), 476–493.
- Dueñas, M. and G. Fagiolo (2013). Modeling the international-trade network: A gravity approach. *Journal of Economic Interaction and Coordination* 8(1), 155–178.
- Erdős, P. and A. Rényi (1959). On random graphs I. *Publicationes Mathematicae Debrecen* 6, 290.
- Fagiolo, G. (2010). The international-trade network: Gravity equations and topological properties. *Journal of Economic Interaction and Coordination* 5(1), 1–25.
- Feng, L., H. Xu, G. Wu, and W. Zhang (2021). Service trade network structure and its determinants in the belt and road based on the temporal exponential random graph model. *Pacific Economic Review* 26(5), 617–650.
- Flandreau, M. and C. Jobst (2005). The ties that divide: A network analysis of the international monetary system, 1890–1910. *Journal of Economic History* 65(4), 977–1007.
- Flandreau, M. and C. Jobst (2009). The empirics of international currencies: Network externalities, history and persistence. *The Economic Journal* 119(537), 643–664.
- Flaschel, P., M. Charpe, G. Galanis, C. R. Proaño, and R. Veneziani (2018). Macroeconomic and stock market interactions with endogenous aggregate sentiment dynamics. *Journal of Economic Dynamics and Control* 91, 237–256.
- Flaschel, P., R. Franke, and W. Semmler (1997). *Dynamic macroeconomics: Instability, fluctuation, and growth in monetary economies*. MIT Press.
- Flaschel, P., G. Groh, C. Proaño, and W. Semmler (2008). *Topics in applied macrodynamic theory*, Volume 10. Springer Science & Business Media.
- Flaschel, P., G. Kauermann, and W. Semmler (2007). Testing wage and price Phillips curves for the United States. *Metroeconomica* 58(4), 550–581.
- Frank, O. and D. Strauss (1986). Markov graphs. *Journal of the American Statistical Association* 81(395), 832–842.
- Fritz, C., G. De Nicola, S. Kevork, D. Harhoff, and G. Kauermann (2023). Modelling the large and dynamically growing bipartite network of German patents and inventors. *Journal of the Royal Statistical Society, Series A: Statistics in Society* 186(3), 557–576.
- Fritz, C., M. Lebacher, and G. Kauermann (2020). Tempus volat, hora fugit: A survey of tie-oriented dynamic network models in discrete and continuous time. *Statistica Neerlandica* 74(3), 275–299.
- Fritz, C., M. Mehrl, P. W. Thurner, and G. Kauermann (2022). Exponential random graph models for dynamic signed networks: An application to international relations. *arXiv preprint*. arXiv:2205.13411.
- Fritz, C., P. W. Thurner, and G. Kauermann (2021). Separable and semiparametric network-based counting processes applied to the international combat aircraft trades. *Network Science* 9(3), 291–311.

- Geyer, C. J. and E. A. Thompson (1992). Constrained Monte Carlo maximum likelihood for dependent data. *Journal of the Royal Statistical Society: Series B (Methodological)* 54(3), 657–683.
- Ghani, A. C., J. Swinton, and G. P. Garnett (1997). The role of sexual partnership networks in the epidemiology of gonorrhea. *Sexually Transmitted Diseases* 24(1), 45–56.
- Gilbert, E. N. (1959). Random graphs. *Annals of Mathematical Statistics* 30(4), 1141–1144.
- Goldenberg, A., A. X. Zheng, S. E. Fienberg, and E. M. Airoldi (2010). A survey of statistical network models. *Foundations and Trends in Machine Learning* 2(2), 129–233.
- Graham, B. S. (2015). Methods of identification in social networks. *Annu. Rev. Econ.* 7(1), 465–485.
- Handcock, M. (2003). Assessing degeneracy in statistical models of social networks. *Working Paper no. 39, University of Washington*.
- Handcock, M. S., D. R. Hunter, C. T. Butts, S. M. Goodreau, P. N. Krivitsky, M. Morris, M. S. Handcock, C. T. Butts, S. M. Goodreau, M. Morris, D. R. Hunter, C. T. Butts, S. M. Goodreau, P. N. Krivitsky, M. Morris, M. S. Handcock, C. T. Butts, S. M. Goodreau, M. Morris, and Martina (2008). ergm: Fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software* 24(3), nihpa54860.
- Handcock, M. S., D. R. Hunter, C. T. Butts, S. M. Goodreau, and M. Morris (2008). statnet: Software tools for the representation, visualization, analysis and simulation of network data. *Journal of statistical software* 24(1), 1548.
- Hanneke, S., W. Fu, and E. P. Xing (2010). Discrete temporal models of social networks. *Electronic Journal of Statistics* 4, 585–605.
- He, X., Y. Dong, Y. Wu, G. Jiang, and Y. Zheng (2019). Factors affecting evolution of the interprovincial technology patent trade networks in china based on exponential random graph models. *Physica A: Statistical Mechanics and its Applications* 514, 443–457.
- Heider, F. (1946). Attitudes and cognitive organization. *Journal of Psychology* 21(1), 107–112.
- Herman, P. R. (2022). Modeling complex network patterns in international trade. *Review of World Economics* 158, 127–79.
- Hoff, P. (2008). Modeling homophily and stochastic equivalence in symmetric relational data. *Advances in Neural Information Processing Systems* 20, 657–664.
- Hoff, P. (2021). Additive and multiplicative effects network models. *Statistical Science* 36(1), 34–50.
- Hoff, P. D. (2005). Bilinear mixed-effects models for dyadic data. *Journal of the American Statistical Association* 100(469), 286–295.

- Hoff, P. D. (2011). Hierarchical multilinear models for multiway data. *Computational Statistics & Data Analysis* 55(1), 530–543.
- Hoff, P. D. (2015). Dyadic data analysis with amen. *arXiv preprint arXiv:1506.08237*.
- Hoff, P. D., A. E. Raftery, and M. S. Handcock (2002). Latent space approaches to social network analysis. *Journal of the American Statistical Association* 97(460), 1090–1098.
- Holland, P. W., K. B. Laskey, and S. Leinhardt (1983). Stochastic blockmodels: First steps. *Social networks* 5(2), 109–137.
- Holland, P. W. and S. Leinhardt (1981). An exponential family of probability distributions for directed graphs. *Journal of the American Statistical Association* 76(373), 33–50.
- Hu, Y. (2005). Efficient, high-quality force-directed graph drawing. *Mathematica Journal* 10(1), 37–71.
- Hummel, R. M., D. R. Hunter, and M. S. Handcock (2012). Improving simulation-based algorithms for fitting ERGMs. *Journal of Computational and Graphical Statistics* 21(4), 920–939.
- Hunter, D. R., S. M. Goodreau, and M. S. Handcock (2008). Goodness of fit of social network models. *Journal of the American Statistical Association* 103(481), 248–258.
- Hunter, D. R. and M. S. Handcock (2006). Inference in curved exponential family models for networks. *Journal of Computational and Graphical Statistics* 15(3), 565–583.
- Hunter, D. R., P. N. Krivitsky, and M. Schweinberger (2012). Computational statistical methods for social network models. *Journal of Computational and Graphical Statistics* 21(4), 856–882.
- Jackson, M. O. (2008). *Social and economic networks*. Princeton University Press.
- Jackson, M. O. (2014). The past and future of network analysis in economics. In *The Oxford Handbook of the Economics of Networks*. Oxford University Press.
- Jackson, M. O. and S. Nei (2015). Networks of military alliances, wars, and international trade. *Proceedings of the National Academy of Sciences* 112(50), 15277–15284.
- Jackson, M. O. and B. W. Rogers (2007). Meeting strangers and friends of friends: How random are social networks? *American Economic Review* 97(3), 890–915.
- Jackson, M. O., B. W. Rogers, and Y. Zenou (2017). The economic consequences of social-network structure. *Journal of Economic Literature* 55(1), 49–95.
- Kindleberger, C. P. (1967). *The politics of international money and world language*. International Finance Section, Department of Economics, Princeton University.
- Koster, J. (2018). Family ties: The multilevel effects of households and kinship on the networks of individuals. *Royal Society Open Science* 5(4), 172159.
- Krivitsky, P. N. (2012). Exponential-family random graph models for valued networks. *Electronic journal of statistics* 6, 1100–1128.

- Krivitsky, P. N. and M. S. Handcock (2014). A separable model for dynamic networks. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 76(1), 29–46.
- Krivitsky, P. N., M. S. Handcock, A. E. Raftery, and P. D. Hoff (2009). Representing degree distributions, clustering, and homophily in social networks with latent cluster random effects models. *Social networks* 31(3), 204–213.
- König, M. D., D. Rohner, M. Thoenig, and F. Zilibotti (2017). Networks in conflict: Theory and evidence from the great war of Africa. *Econometrica* 85(4), 1093–1132.
- Lebacher, M., P. W. Thurner, and G. Kauermann (2021). A dynamic separable network model with actor heterogeneity: An application to global weapons transfers. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 184(1), 201–226.
- Lee, Y. and E. L. Ogburn (2021). Network dependence can lead to spurious associations and invalid inference. *Journal of the American Statistical Association* 116(535), 1060–1074.
- Lewer, J. J. and H. Van den Berg (2008). A gravity model of immigration. *Economics Letters* 99(1), 164–167.
- Liu, L., M. Shen, D. Sun, X. Yan, and S. Hu (2022). Preferential attachment, R&D expenditure and the evolution of international trade networks from the perspective of complex networks. *Physica A: Statistical Mechanics and its Applications* 603, 127579.
- Lusher, D., J. Koskinen, and G. Robins (2012). *Exponential Random Graph Models for Social Networks*. Cambridge University Press.
- Matias, C. and S. Robin (2014). Modeling heterogeneity in random graphs through latent space models: A selective review. *ESAIM: Proceedings and Surveys* 47, 55–74.
- Mele, A. (2017). A structural model of dense network formation. *Econometrica* 85(3), 825–850.
- Miller, S. V. (2022). {peacesciencer}: An R package for quantitative peace science research. *Conflict Management and Peace Science* 39(6), 755–779.
- Minhas, S., C. Dorff, M. B. Gallop, M. Foster, H. Liu, J. Tellez, and M. D. Ward (2022). Taking dyads seriously. *Political Science Research and Methods* 10(4), 703–721.
- Minhas, S., P. D. Hoff, and M. D. Ward (2016). A new approach to analyzing coevolving longitudinal networks in international relations. *Journal of Peace Research* 53(3), 491–505.
- Minhas, S., P. D. Hoff, and M. D. Ward (2019). Inferential approaches for network analysis: AMEN for latent factor models. *Political Analysis* 27(2), 208–222.
- Morales, E., G. Sheu, and A. Zahler (2019). Extended gravity. *Review of Economic Studies* 86(6), 2668–2712.
- Morris, M., M. S. Handcock, and D. R. Hunter (2008). Specification of exponential-family random graph models: Terms and computational aspects. *Journal of Statistical Software* 24(4), 1548.

- Mundt, P. (2021). The formation of input–output architecture: Evidence from the European Union. *Journal of Economic Behavior & Organization* 183, 89–104.
- Nelder, J. A. and R. W. M. Wedderburn (1972). Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)* 135(3), 370–384.
- Newcomb, T. M. (1979). Reciprocity of interpersonal attraction: A nonconfirmation of a plausible hypothesis. *Social Psychology Quarterly* 42(4), 299–306.
- Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review* 45(2), 167–256.
- R Core Team (2021). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing.
- Rivera, M. T., S. B. Soderstrom, and B. Uzzi (2010). Dynamics of dyads in social networks: Assortative, relational, and proximity mechanisms. *Annual Review of Sociology* 36, 91–115.
- Robins, G., P. Pattison, Y. Kalish, and D. Lusher (2007). An introduction to exponential random graph (p^*) models for social networks. *Social Networks* 29(2), 173–191.
- Robins, G., T. Snijders, P. Wang, M. Handcock, and P. Pattison (2007). Recent developments in exponential random graph (p^*) models for social networks. *Social Networks* 29(2), 192–215.
- Rose, A. K. (2004). Do we really know that the WTO increases trade? *American Economic Review* 94(1), 98–114.
- Schoeneman, J., B. Zhu, and B. A. Desmarais (2022). Complex dependence in foreign direct investment: Network theory and empirical analysis. *Political Science Research and Methods* 10(2), 243–259.
- Schweinberger, M. (2011). Instability, sensitivity, and degeneracy of discrete exponential families. *Journal of the American Statistical Association* 106(496), 1361–1370.
- SIPRI (2021). SIPRI Arms Transfers Database. *Stockholm International Peace Research Institute*. <https://www.sipri.org/databases/armstransfers> (accessed 10/20/2022).
- Smith, M., S. Gorgoni, and B. Cronin (2019). International production and trade in a high-tech industry: A multilevel network analysis. *Social Networks* 59, 50–60.
- Smith, M. and Y. Sarabi (2022). How does the behaviour of the core differ from the periphery? – An international trade network analysis. *Social Networks* 70, 1–15.
- Snijders, T. A. (1996). Stochastic actor-oriented models for network change. *Journal of mathematical sociology* 21(1-2), 149–172.
- Snijders, T. A. (2017). Stochastic actor-oriented models for network dynamics. *Annual review of statistics and its application* 4, 343–363.
- Snijders, T. A. B., P. E. Pattison, G. L. Robins, and M. S. Handcock (2006). New specifications for exponential random graph models. *Sociological Methodology* 36(1), 99–153.

- Thurner, P. W., C. S. Schmid, S. J. Cranmer, and G. Kauermann (2019). Network interdependencies and the evolution of the international arms trade. *Journal of Conflict Resolution* 63(7), 1736–1764.
- van der Pol, J. (2019). Introduction to network modeling using exponential random graph models (ergm): Theory and an application using R-project. *Computational Economics* 54(3), 845–875.
- Ward, M. D., J. S. Ahlquist, and A. Rozenas (2013). Gravity’s rainbow: A dynamic latent space model for the world trade network. *Network Science* 1(1), 95–118.
- Ward, M. D. and P. D. Hoff (2007). Persistent patterns of international commerce. *Journal of Peace Research* 44(2), 157–175.
- Wasserman, S. and K. Faust (1994). *Social Network Analysis: Methods and Applications*. Cambridge University Press.
- Wasserman, S. and P. Pattison (1996). Logit models and logistic regressions for social networks: I. An introduction to Markov graphs and p^* . *Psychometrika* 61(3), 401–425.

Supplementary Material

Dependence matters: Statistical models to identify the drivers of tie formation in economic networks

Giacomo De Nicola^{1,*}, Cornelius Fritz², Marius Mehrl³, Göran Kauermann¹

Department of Statistics, LMU Munich¹

Department of Statistics, Pennsylvania State University²

School of Politics and International Studies, University of Leeds³

The following contains technical details and supplementary information to the manuscript. The full data and code to reproduce the analysis and to aid the reader in applying the presented models to their own data can be found in our Github repository, available at <https://github.com/gdenicola/statistical-network-analysis-in-economics>.

S.1 Details on the p_1 model

To represent reciprocity, the p_1 model assumes dyads, defined by (Y_{ij}, Y_{ji}) , to be independent of one another. The resulting bivariate distribution of each dyad (Y_{ij}, Y_{ji}) comprises three parameters and a constraint:

$$\begin{aligned} m_{ij} &= \mathbb{P}(Y_{ij} = Y_{ji} = 1) \\ a_{ij} &= \mathbb{P}(Y_{ij} = 1, Y_{ji} = 0) \\ e_{ij} &= \mathbb{P}(Y_{ij} = Y_{ji} = 0) \\ m_{ij} + a_{ij} + a_{ji} + e_{ij} &= 1. \end{aligned}$$

The joint distribution of the network is then given by:

$$\begin{aligned} \mathbb{P}_{\theta}(\mathbf{Y} = \mathbf{y}) &= \prod_{i < j} m_{ij}^{y_{ij}y_{ji}} a_{ij}^{y_{ij}(1-y_{ji})} a_{ji}^{(1-y_{ij})y_{ji}} e_{ij}^{(1-y_{ij})(1-y_{ji})} \\ &= \frac{\exp \left\{ \sum_{i < j} \rho_{ij} y_{ij} y_{ji} + \sum_{i \neq j} \theta_{ij} y_{ij} \right\}}{\kappa(\theta)}, \end{aligned} \tag{1}$$

where $\rho_{ij} = \log(m_{ij}e_{ij}/a_{ij}^2)$ for $i < j$, $\theta_{ij} = \log(a_{ij}/n_{ij})$ for $i \neq j$, and $\kappa(\theta)$ is the normalizing constant. To estimate the parameters in (1) we need to introduce some homogeneity assumption to avoid overparametrization. Following Holland and Leinhardt (1981), we assume

$$\begin{aligned} \rho_{ij} &= \theta_{\text{Rep}} \quad \forall i < j \\ \theta_{ij} &= \theta_{\text{Edges}} + \theta_{\text{Out},i} + \theta_{\text{In},j} \quad \forall i \neq j \\ \sum_{i=1}^n \theta_{\text{Out},i} &= \sum_{i=1}^n \theta_{\text{In},i} = 0, \end{aligned}$$

where θ_{Rep} is the global effect of reciprocity, θ_{Edges} quantifies the general sparsity in the network, and $\theta_{\text{Out},i}$ and $\theta_{\text{In},i}$ indicate the general tendency for all agents $i \in \{1, \dots, n\}$ to

*Corresponding author: giacomo.denicola@stat.uni-muenchen.de

form out- or in-going ties. Combining these homogeneity assumptions with (1) yields

$$\mathbb{P}_{\theta}(\mathbf{Y} = \mathbf{y}) \propto \exp \left\{ \theta_{\text{Rep}} s_{\text{Rep}}(\mathbf{y}) + \theta_{\text{Edges}} s_{\text{Edges}}(\mathbf{y}) + \sum_{i=1}^n \theta_{\text{Out},i} s_{\text{Out},i}(\mathbf{y}) + \sum_{i=1}^n \theta_{\text{In},i} s_{\text{In},i}(\mathbf{y}) \right\}, \quad (2)$$

where $s_{\text{Rep}}(\mathbf{y}) = \sum_{i < j} y_{ij} y_{ji}$ is the number of reciprocal ties, $s_{\text{Edges}}(\mathbf{y}) = \sum_{i=1}^n \sum_{j \neq i} y_{ij}$ the number of edges, and $s_{\text{Out},i}(\mathbf{y})$ and $s_{\text{In},i}(\mathbf{y})$ the number of out- and in-going ties of agent i .

Note that all these statistics are functions of the observed network and are the sufficient statistics, i.e. they contain all necessary information for determining all coefficients in (2). This sufficiency principle translates to $s_{\text{Out},i}(\mathbf{y}) = s_{\text{Out},j}(\mathbf{y}) \Rightarrow \theta_{\text{Out},i} = \theta_{\text{Out},j}$. This, in turn, allows us to write all agent-specific terms in (2) in terms of degree statistics:

$$\mathbb{P}_{\theta}(\mathbf{Y} = \mathbf{y}) \propto \exp \left\{ \theta_{\text{Rep}} s_{\text{Rep}}(\mathbf{y}) + \theta_{\text{Edges}} s_{\text{Edges}}(\mathbf{y}) + \sum_{i=1}^n \theta_{\text{Outdeg},i} s_{\text{Outdeg},i}(\mathbf{y}) + \sum_{i=1}^n \theta_{\text{Indeg},i} s_{\text{Indeg},i}(\mathbf{y}) \right\}, \quad (3)$$

where $s_{\text{Outdeg},i}(\mathbf{y})$ and $s_{\text{Indeg},i}(\mathbf{y})$ are statistics counting the actors with out/in-degree i in $\mathbf{y}$.

S.2 Further results of the application to the international arms trade network

S.2.1 Model diagnostics and goodness of fit

To evaluate the goodness of fit of an estimated ERGM, the standard approach is to compare the statistics observed in the real world network with the distribution of the same statistics calculated on networks simulated from the model (Hunter et al., 2008). The results of this comparison are depicted in Figure S.1, where we investigate the goodness of fit of the ERGM estimated on the 2018 international arms trade network. The black line in each subfigure represents the distribution of the respective network statistic observed in the real network. For instance, the top right figure indicates that approximately 10% of all nodes (countries) in the network had an in-degree of 2. The boxplots then show the distribution of each value of a given statistic over the simulated networks. A good model will thus generally result in boxplots that include the observed values of the network statistics under consideration. In Figure S.1, this is almost always the case, though it is also visible that the real network included less countries with an in-degree of 0 but more with an in-degree of 1 than the large majority of networks simulated from the ERGM.

In Figure S.2, we further compare the performance of the fully specified ERGM against that of the logistic regression model including the same exogenous covariates as the ERGM, but of course omitting all endogenous network statistics. We already noted that the ERGM appears to do a better job at in-sample prediction than the Logit model given its lower AIC and BIC values. Figure S.2 documents both models' respective areas under the receiver-operator (ROC) and precision-recall curves (PR). Here, a higher value of each curve indicates better predictive performance, and again, the ERGM appears to outperform the logistic regression model for both metrics. Figure S.2 thus offers further evidence that in the case of the international arms trade, model performance is improved by accounting for endogenous network effects.




Figure S.1: Goodness of fit plots for the application to the international arms trade network in 2018. The metrics show a reasonably good fit for our model.




Figure S.2: In-sample prediction performance of the two models in Table 1. Left panel: ROC. Right panel: PR curve. ERGM outperforms Logit in both metrics.

S.2.2 Robustness checks

In the analysis presented in Section 3.3, we lag all covariate information by three years. Since this choice is based on the heuristic that there are considerable delays between the order and delivery date of arms, we next show the results of the ERGM with different time lags, namely four and five years, in Table S.1. From these results, we conclude that the results are robust to different specifications of the time lags between the modeled network and the exogenous covariates. Moreover, one can observe that the ERGM consistently achieves lower AIC and BIC values compared to its classical logistic counterpart, underpinning the finding that accounting for network dependencies is necessary for analyzing the international arms trade. We further corroborate this point through a ROC as well as a PR analysis, both shown in Figures S.3(a) and (b).

Covariate Lag	ERGM 4 years	Logit 4 years	ERGM 5 years	Logit 5 years
Edges	−15.46 (2.10)	−28.22 (1.73)	−13.82 (3.19)	−23.11 (3.04)
Repetition	3.25 (0.14)	3.96 (0.14)	3.25 (0.14)	3.95 (0.14)
Distance	−0.08 (0.09)	−0.24 (0.09)	−1.68 (2.49)	−5.93 (2.47)
Abs. Diff. Polity	−0.00 (0.01)	−0.00 (0.01)	−0.00 (0.01)	−0.00 (0.01)
Alliance	0.37 (0.20)	0.21 (0.21)	0.40 (0.20)	0.24 (0.21)
log-GDP (Sender)	0.30 (0.05)	0.59 (0.04)	0.30 (0.05)	0.59 (0.04)
log-GDP (Receiver)	0.17 (0.05)	0.35 (0.04)	0.16 (0.05)	0.35 (0.04)
Mutual	−0.34 (0.43)		−0.33 (0.44)	
GWIDEG	−1.48 (0.31)		−1.49 (0.30)	
GWODEG	−2.84 (0.30)		−2.85 (0.30)	
GWOTP	−0.15 (0.11)		−0.15 (0.11)	
GWISP	0.22 (0.09)		0.22 (0.09)	
AIC	1772.34	1891.65	1773.36	1893.14
BIC	1868.67	1947.84	1869.69	1949.34
Log Likelihood	−874.17	−938.82	−874.68	−939.57

Table S.1: Estimated coefficients of ERGMs and corresponding logistic regression models for the international arms trade network in 2018 with covariates lagged by 4 and 5 years, respectively.




Figure S.3: In-sample prediction performance of the four models in Table S.1. Upper row: Results when lagging the covariates four years. Lower row: Results when lagging the covariates five years. Left column: ROC. Right column: PR curve. ERGM outperforms Logit in both cases and for both metrics.

S.3 Further results of the application to the global foreign exchange activity network

S.3.1 Variance and covariance parameters

In Section 4.3 of the paper, we showcased the AME model by fitting it to the historical network of global foreign exchange activity in 1900. In illustrating the results we, for brevity, focused on the main effects of the covariates included, reported in Table 2. But the AME model also estimates several variance and covariance parameters, which also have a meaningful interpretation. The estimates for those parameters are reported in Table S.2. The first two parameters, σ_a^2 and σ_b^2 , represent the estimated variances of

	Estimate	Standard deviation
σ_a^2	0.557	0.165
σ_b^2	1.863	0.860
σ_{ab}	-0.350	0.250
ρ	0.364	0.319

Table S.2: Estimated variance and covariance parameters for the AMEN model fitted on the global foreign exchange network in 1900.

the additive sender and receiver effects, respectively. We can see that receiver effects are much more variable than sender effects. This makes intuitive sense given that there are a few currencies which are traded by a large number of countries, while many currencies aren't traded at all outside of their origin countries. The skewness in the distribution of incoming ties thus induces a relatively large variance in the receiver effects.

The coefficient in the third row, σ_{ab} , measures the (global) correlation between sender and receiver effects of the same node. In this case, we can see that there is a slight negative correlation between the two, meaning that countries that trade many foreign currencies within their financial hub do not necessarily tend to have their home currency traded in many countries. Finally, the coefficient ρ indicates the covariance between the residuals on the same node-pair, ϵ_{ij} and ϵ_{ji} . This parameter quantifies the tendency towards reciprocity in the network. In this case, we can see that there is a slight positive tendency for ties to be reciprocated.

S.3.2 Model diagnostics and goodness of fit

Similarly as for ERGM, the goodness of fit for AME models is evaluated by comparing the statistics observed in the real world network with the distribution of the same statistics calculated for networks simulated from the model. Figure S.4 depicts this comparison for first order effects (top panel) and for second- and third-order dependencies (bottom panel), as done by default in the R package `amen` (see Hoff, 2015 for details). All in all, we can see that the model does a reasonably good job in preserving the network statistics in question.

We can also compare the goodness of fit of the AME model with that of the Probit model including the same exogenous covariates, but of course omitting the latent variables. Figure S.5 depicts the same comparison of observed with simulated statistics just described, but for Probit instead of AME. From the plots therein we can see how the Probit model does a markedly worse job than the AME in reproducing first and third order dependencies, thus demonstrating an overall worse performance in capturing the mechanisms at play in the network.

In Figure S.6, also produced by default by the `amen` package, we can further check how the coefficients and their variance vary across the MCMC iterations. In general, the fit is considered to be acceptable if no visible trends emerge in the chains. If, to the contrary, trends in the estimates are visible, the researcher must consider running the chain for more iterations and/or using alternative model formulations. In the case of Figure S.6, visual inspection gives us confidence that the AME model has reasonably converged.




Figure S.4: Goodness of fit plots for the AME model applied to the global foreign exchange network in 1900. The model does an overall good job in replicating the observed network statistics.




Figure S.5: Goodness of fit plots for the classical probit model applied to the global foreign exchange network in 1900. The model performs visibly worse than AME in replicating the observed network statistics.




Figure S.6: MCMC diagnostics for the AME model applied to the global foreign exchange network in 1900. Visual inspection gives us confidence that the model has reasonably converged.