

RESEARCH ARTICLE

Random effects in dynamic network actor models

Alvaro Uzaheta^{1*}, Viviana Amati² and Christoph Stadtfeld¹

¹Social Networks Lab, ETH Zurich, Zurich, Switzerland and ²Department of Statistics and Quantitative Methods, University of Milano-Bicocca, Milan, Italy

*Corresponding author. Email: alvaro.uzaheta@gess.ethz.ch

Guest Editors (Special Issue on Relational Event Models): Carter Butts, Alessandro Lomi, Tom Snijders

Abstract

Dynamic Network Actor Models (DyNAMs) assume that an observed sequence of relational events is the outcome of an actor-oriented decision process consisting of two decision levels. The first level represents the time until an actor initiates the next relational event, modeled by an exponential distribution with an actor-specific activity rate. The second level describes the choice of the receiver of the event, modeled by a conditional multinomial logit model. The DyNAM assumes that the parameters are constant over the actors and the context. This homogeneity assumption, albeit statistically and computationally convenient, is difficult to justify, e.g., in the presence of unobserved differences between actors or contexts. In this paper, we extend DyNAMs by including random-effects parameters that vary across actors or contexts and allow controlling for unknown sources of heterogeneity. We illustrate the model by analyzing relational events among the users of an online community of aspiring and professional digital and graphic designers.

Keywords: Bayesian estimation; Dynamic Network Actor Models; heterogeneity; random-effects models; relational event data

1. Introduction

Social relationships are formed and maintained through individuals' interpersonal actions (Borgatti & Halgin, 2011). These actions are often collected as relational events, i.e. sequences of directed interactions between individuals. A relational event, in general, consists of at least three elements (Stadtfeld & Block, 2017): a social actor who generates an event (the “sender”), a social actor who receives the event (the “receiver”), and the time at which the event takes place. Relational events occur in a variety of contexts, including interpersonal interactions (Stadtfeld & Block, 2017), online peer-production (Lerner & Lomi, 2017), treaties between countries (Stadtfeld et al., 2017a), classroom conversations (DuBois et al., 2013), animal behavior (Tranmer et al., 2015; Patison et al., 2015), online political discussion (Mulder & Leenders, 2019), and capital transfer between banks (Zappa & Vu, 2021), among others.

Several statistical models for analyzing relational event data have been proposed (Butts, 2008; Perry & Wolfe, 2013; Stadtfeld, 2012; Vu et al., 2011a). These models enable the testing of hypotheses about how previous events, nodal attributes, dyadic attributes, and network structures explain the emergence of relational events through time. Among the models, Dynamic Network Actor Models (DyNAMs) investigate relational events as an actor-oriented decision process consisting of two decision levels (Stadtfeld & Block, 2017). The first level represents the time until an actor initiates the next relational event. The time between two consecutive events originated by the same actor is modeled as an exponential distribution with an actor-specific activity rate. The

second level describes the choice of the receiver of the event. A conditional multinomial logit model determines the probability that an actor is chosen as a receiver.

Combinations of statistics and parameters define both the actor activity rate and the multinomial linear predictor. The statistics of the rates explain differences in actors' activity due to their position in the network, their characteristics, or the sequence of previous events. The statistics of the multinomial logit model express preferences for various types of social structures, such as reciprocated dyads, transitive triads, or homophilous dyads. These preferences may vary depending on actors' characteristics and their current position in the network. The parameters quantify the effect of the statistics in the evolution of the process. The DyNAM assumes that the parameters are constant over the actors and the context. Thus, heterogeneity between actors enters the model through statistics accounting for differences in actors' attributes and structural positions that can be observed. In other words, the model assumes that actors with the same characteristics and network positions follow the same behavioral patterns. This homogeneity assumption, albeit statistically and computationally convenient, is difficult to justify in some empirical settings, e.g., in the presence of unobserved differences between actors or contexts. Such differences can lead to unobserved heterogeneity in actors' activity or preferences beyond what the model specification could explain.

In this paper, we extend DyNAMs by including random-effects parameters that vary across actors or contexts and allow controlling for unknown sources of heterogeneity. The extension is rooted in the literature of multilevel and mixed-effects models (Gelman & Hill, 2006; Snijders & Bosker, 2011; Raudenbush & Bryk, 2002). Such models are commonly applied when the data presents a nested structure (e.g., students in classrooms and longitudinal measurement of subjects). The multilevel approach is suitable since we can think of the relational event sequence as a series of events nested within the sending actors or a pre-specified partition of the actors in groups. We follow a Bayesian approach to make inferences about the model parameters. The procedure uses Hamiltonian Chain Monte Carlo with the No-U-turn sampler (NUTS) as implemented in Stan (Hoffman & Gelman, 2014; Stan Development Team, 2022).

The model developed in this study is new in two regards. First, the model allows explaining actors' (or groups') deviation from the average tendency of a given network mechanism using actors' covariates or monadic network statistics, thereby controlling for the differential agency of actors. Second, it introduces random effects to account for and test hypotheses on unobserved heterogeneity between actors for specific network statistics.

The article is organized as follows. Section 2 summarizes the literature on the relational event and random-effects network models. Section 3 presents the logic and the formulation of the standard DyNAM. The definition of the random-effects DyNAMs is described in Section 4. Sections 5 and 6 report information on the model estimation and selection. Section 7 illustrates the model by analyzing relational events among an online community of aspiring and professional digital and graphic designers. The paper concludes in Section 8 with a discussion.

2. Review of existing statistical network models

The DyNAM with random effects leans on two research lines: the Relational Event Model (REM) and network models that include random effects. Butts (2008) introduced the REM, a model inspired by event history analysis and agent-based modeling. This model assumes that all possible sender–receiver pairs are competing to initiate an event at each point in time. Under the assumption that the events occur following an inhomogeneous Poisson process and are independent conditional on the realized history of the process, it is possible to link the occurrence of the events with a rate function. This function depends on the information of the particular dyad and its position in the events' network.

In contrast, the DyNAM uses ideas similar to the Stochastic Actor-Oriented Model (Snijders, 1996, 2001). It proposes a two-step decision process embedded in a continuous-time Markov chain. The first step describes the time until an actor initiates the next relational event and is modeled by an exponential distribution depending on the actor activity rate. The second determines the choice of the receiver of the event and is modeled by a multinomial probability distribution.

Variations describing the social interactions as the outcome of point processes have been presented by Perry and Wolfe (2013) and Vu et al. (2011a).

The inclusion of random effects in network models has been proposed in the literature for both cross-sectional and longitudinal network data. The p_2 model (Van Duijn et al., 2004), Exponential Random Graph Models with sender and receiver random effects (Box-Steffensmeier et al., 2018), and with random effects for a given network statistic (Thiemichen et al., 2016) have been developed to account for actor heterogeneity in the analysis of cross-sectional data. Different work lines close to our proposal led to the extensions of the Stochastic Actor-Oriented Model and REM for longitudinal data. For the Stochastic Actor-Oriented Model, Schweinberger (2007, chap. 4) and Schweinberger (2020) proposed random effects to account for variations between actors and groups of actors, respectively. For the REM, DuBois et al. (2013) studied multiple event sequences introducing random effects that vary across contexts.

3. Dynamic Network Actor Models

Following the notation of Butts (2008) and Stadtfeld and Block (2017), we represent an event as a tuple of the form $\omega = (i_\omega, j_\omega, t_\omega)$, where $i_\omega \in \mathcal{S}(t_\omega)$ denotes the event sender, $j_\omega \in \mathcal{R}(t_\omega)$ is the event receiver, and $t_\omega \in \mathbb{R}_{\geq 0}$ represents the time at which the event occurred. The sets $\mathcal{S}(t_\omega)$ and $\mathcal{R}(t_\omega)$ represent the collection of individuals, entities or agents that could potentially be the senders or receivers of an event at time t_ω . These sets can be disjoint ($\mathcal{S}(t_\omega) \cap \mathcal{R}(t_\omega) = \emptyset$), as in the case of two-mode (bipartite) relational event data, or identical ($\mathcal{S}(t_\omega) = \mathcal{R}(t_\omega)$), as in the case of one-mode relational event data.

A sequence of time-ordered relational events $\Omega = \{\omega_1, \omega_2, \dots, \omega_N\}$ is the focus of the model as the dependent variable. The process state $\mathbf{y}(t)$ is the explanatory variable in the model. It comprises all the relevant information of past events and attributes up to but not including t_ω . Formally, it is a multivariate system defined as

$$\mathbf{y}(t) = (\mathbf{x}^{(1)}(t), \mathbf{x}^{(2)}(t), \dots, \mathbf{z}^{(1)}(t), \mathbf{z}^{(2)}(t), \dots, \mathcal{S}(t), \mathcal{R}(t)) \quad t < t_\omega. \tag{1}$$

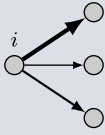
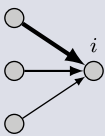

The elements $\mathbf{x}^{(i)}(t)$ are matrices representing the sequence of previous events, various types of networks, or relational data. The elements $\mathbf{z}^{(i)}(t)$ are nodal or global covariates that can be expressed as vectors or scalars. We omit the time indicator t when convenient, and denote vectors and matrices by bold case letters.

DyNAMs (Stadtfeld & Block, 2017) assume that the observed sequence of relational events is the outcome of a continuous-time Markov process with two sub-steps modeling the changes in the network. The first step describes the time until an actor initiates the next relational event. The waiting time of an actor i is modeled by an exponential distribution with rate parameter τ_i that depends on a parameter vector $\boldsymbol{\theta}$, and an arbitrary set of real functions r_g evaluating the process state:

$$\tau_i(\mathbf{y}; \boldsymbol{\theta}) = \exp\left(\theta_0 + \sum_g \theta_g r_g(i, \mathbf{y})\right). \tag{2}$$

The functions r_g , hereafter referred to as statistics, convey the idea that actors might exhibit a different rate due to their network position or covariates. The parameter θ_g represents the effect of the statistic r_g in determining the rate. Positive values of the parameter indicate that the rate increases with the increase of r_g , thereby implying that actors with high values of r_g are more likely

Table 1. Examples of statistics r_g for the rate function τ_i in equation (2)

Statistic r_g	Representation	Interpretation of θ_g
Out-degree activity		Tendency of actor i to initiate an event when i has a high outgoing degree
In-degree popularity		when i has a high incoming degree
Sender attribute		i scores high on attribute

The interpretation refers to positive values of θ_g . A similar interpretation holds for negative values of θ_g .

to initiate events than actors with low values of r_g . Table 1 provides the pictorial representation of a few statistics r_g along with the interpretation of the corresponding parameters θ_g .

The second step describes the choice of the receiver of the event. According to random utility theory, the choice of the receiver is modeled by a multinomial model having the form:

$$p(i \rightarrow j; \mathbf{y}, \boldsymbol{\beta}) = \frac{\exp(f(i, j, \mathbf{y}, \boldsymbol{\beta}))}{\sum_k \exp(f(i, k, \mathbf{y}, \boldsymbol{\beta}))}, \forall j. \tag{3}$$

In equation (3), $f(i, j, \mathbf{y}, \boldsymbol{\beta})$ represents the systematic part of the utility function that actor i uses to evaluate all possible alters. The systematic component is a linear combination of real functions s_h , evaluating the process state in the presence of the event $i \rightarrow j$ and parameters β_h :

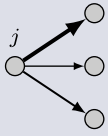
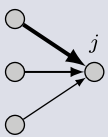



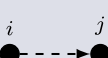
$$f(i, j, \mathbf{y}, \boldsymbol{\beta}) = \sum_h \beta_h s_h(i, j, \mathbf{y}). \tag{4}$$

The functions s_h are statistics expressing the sender preferences for actors having specific characteristics or embedded in particular social structures. The parameter β_h represents the effect of the statistic s_h in determining the receiver. Positive values of the parameter indicate that events to receiver j leading to high values of s_h are more likely than those to receiver k leading to low values of s_h . Table 2 shows the pictorial representation of the main statistics s_h along with the interpretation of the corresponding parameters β_h .

The subprocesses describing the waiting time between events and the choice of the receiver are assumed to be conditionally independent, given the process state (Stadtfeld & Block, 2017). Thus, the vector of parameters $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$ can be obtained by performing estimation using the partial likelihoods of the submodels for the actor activity rates and the multinomial receiver choices, respectively. The R package `goldfish` (Stadtfeld & Hollway, 2022) provides the implementation of DyNAMs.

It must be noted that the assumption of separability in DyNAMs is analytically convenient, but also a reasonable representation of many social processes under study. It may not be tenable in situations where the timing of events depends on the choice of the receiver. We refer the readers

Table 2. Examples of statistics s_h for the systematic component $f(i, j, \mathbf{y}, \boldsymbol{\beta})$ in equation (3).

Statistic s_h	Representation	Interpretation of β_h
Out-degree activity		Tendency of creating an event $i \rightarrow j$ when j has a high outgoing degree
In-degree Popularity		when j has a high incoming degree
Reciprocity		if one or several events $j \rightarrow i$ happened in the past
Inertia		if one or several events $i \rightarrow j$ happened in the past
Receiver attribute		when j scores high on attribute
Homophily		if actors i and j have the same attribute

The interpretation refers to positive values of β_h . A similar interpretation holds for negative values of β_h .

to the papers of Stadtfeld and Block (2017) and Stadtfeld et al. (2017b) for an extensive discussion on the assumption of separability.

4. Random effects in DyNAMs

The parameters θ_g and β_h in equations (2) and (4), respectively, are assumed to be constant over the actors, thereby implying that the effect of the statistics is the same for all actors with the same characteristics and network positions. Here, we present the mathematical formulation of the random effects model for the second sub-step (the receiver’s choice) to relax the assumption that all actors have the same preferences for events leading to the formation of the statistics s_h . An analogous extension can be derived for the rate parameter τ_i .

The derivation of the probability of choosing a receiver as a multinomial choice model under a random utility maximization is based on the assumptions of independence of irrelevant alternatives and the presence of additive disturbances distributed as a Type I extreme value distribution (McFadden, 1974; Luce & Suppes, 1965). Additionally, the assumption that the random disturbances of the utility function are independent and identically distributed implies that the error term is a noise arising from unknown and unspecified sources after controlling for the covariates included in the model. Translating this into the framework of DyNAMs, this assumption implies that the receiver choices made from the same actor are independent after controlling for the state

of the process through the statistics s_h . However, consecutive choices made by the same actor present correlation even after controlling for covariates or could vary for the presence of random taste variations.

This problem has been recognized in the discrete choice modeling literature. McFadden and Train (2000) developed the random-coefficients logit model where the multinomial coefficients can vary over individuals and alternatives.

We use the same idea to extend the formulation of the utility function to include random effects in DyNAMs. The model specification follows the notation of Gelman and Hill (2006). The utility function is split into two components as represented by the following equation

$$f(i, j, \mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\alpha}) = \sum_h \beta_h s_h(i, j, \mathbf{y}) + \sum_m \alpha_{mi} s_m(i, j, \mathbf{y}). \quad (5)$$

The first component resembles equation (4) and represents the local configurations that are evaluated in the same way by all the actors as described by the fixed-effect coefficients β_h . The second part introduces a new set of statistics s_m that depends on a set of random-effect coefficients α_{mi} for actor i . The parameters α_{mi} are random slopes that allow the fixed effect part to vary over the actors. Formally, α_{mi} is defined as:

$$\alpha_{mi} = \gamma_{m0} + \sum_n \gamma_{mn} a_n(i, \mathbf{y}) + \eta_{mi}$$

$$\boldsymbol{\eta}_i \sim MN(\mathbf{0}, \boldsymbol{\Sigma}).$$

The parameter γ_{m0} represents the average importance of statistic s_m in the population when all real functions a_n take value zero. The second component $\sum_n \gamma_{mn} a_n(i, \mathbf{y})$ controls for variations of the importance of statistic s_m that are explained by the real function a_n evaluating the process state. The functions a_n can be nodal covariates or network statistics that reflect differences in the network position of the actors, similarly to the statistics r_g used in the waiting time formulation. It is worth mentioning that this second component is not considered in previous formulations of the random-coefficients logit model introduced in McFadden and Train (2000).

The third component η_{mi} captures the variations from the mean population importance γ_{m0} that is not explained by the actor's statistics a_n . Here, we assume that when the number of random effects is larger than two, the $\boldsymbol{\eta}_i$ has a multivariate normal distribution with mean vector $\mathbf{0}$ and variance-covariance matrix $\boldsymbol{\Sigma}$. In the case of one random effect, this assumption reduces to a univariate normal distribution with mean 0 and variance σ^2 .

The formulation of the random-effects DyNAM presented above is quite flexible. The introduction of random effects controlling for unobserved heterogeneity in actors' activity (i.e. in the rate function τ_i) is analogous to that in the receiver's choice. Moreover, the use of random effects in DyNAM can be extended to model unobserved heterogeneity in data with different nested structures, in analogy with the modeling of multiple sequences introduced by DuBois *et al.* (2013) for REMs. Compared to the model of DuBois *et al.* (2013), the random-effects DyNAM is more general since it includes nodal covariates a_n in the definition of the random slopes.

The proposed model is implemented in the R package `goldfish.latent` available and hosted on GitHub <https://github.com/snlab-ch/goldfish.latent>.

5. Model estimation

The assumption that the subprocesses describing the waiting time between events and the choice of the receiver are conditionally independent, given the process state, holds in the presence of random effects. Thus, the vectors of parameters $\boldsymbol{\theta}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$ are estimated separately when only

one of the submodels includes random effects. If the random effects appear in both submodels, the separability assumption does not hold anymore, and the parameters need to be estimated jointly.

Here, we describe the estimation of a DyNAM with random effects in the receiver choice model only. The estimation of the parameter θ is carried out over the partial likelihood as described at the end of Section 3.

We follow a Bayesian approach to make inferences about the random-effect as well as the fixed-effect coefficients in the receiver choice submodel. The parameters are estimated using the Hamiltonian Monte Carlo (HMC) method (Neal, 2011; Hoffman & Gelman, 2014) as implemented in the probabilistic programming language Stan (Stan Development Team, 2022). The HMC method and model parametrization we chose allows circumventing two issues that usually happen when estimating the parameters of mixed-effect models.

The first issue concerns the number of parameters to be estimated. It is well-known that inference for high-dimensional models is challenging. Typical Markov Chain Monte Carlo algorithms, like Gibbs or Metropolis sampling, tend to get stuck in local neighborhoods and poorly explore the space of the posterior distribution due to the use of inefficient random walks. It follows that sampling from the posterior distribution of mixed-effect models might be extremely time-consuming.

Algorithms based on the HMC have been developed to overcome this issue (Neal, 2011). More specifically, the HMC method transforms the problem of sampling from a target distribution into the problem of simulating Hamiltonian dynamics by introducing an auxiliary variable scheme. In a nutshell, HMC explores the sample space simulating the evolution of a particle in a Hamiltonian system where the position and momentum of the particle are updated step by step. Empirical studies have proven that HMC algorithms perform much better than the Gibbs and Metropolis algorithms in high-dimensional scenarios. We refer the readers interested in the intuition behind the HMC method and the technical details to the papers of Betancourt (2017) and Neal (2011).

The second issue relates to the parametrization of the distribution for the random effects. As described in Betancourt and Girolami (2015), small changes in the parameters of the random effects distributions might induce considerable changes in the density of the posterior distribution. If the data are sparse, the posterior distribution has the shape of a funnel with a quite narrow neck. A method to circumvent this pathology is to reparametrize the model with the help of auxiliary variables in such a way that the search space consists of independent variables. Papaspiliopoulos et al. (2007) introduced the notion of scale noncentered parametrization for the mixed-effect model.

For the random-effects DyNAM, the scale noncentered parametrization is obtained by factorizing the variance-covariance matrix Σ using the Cholesky factorization. The factorization represents the random effects as a transformation of a set of independent normal standard variables:

$$\begin{aligned}\alpha_{mi} &= \gamma_{m0} + \sum_n \gamma_{mn} a_n(i, \mathbf{y}) + \tilde{\eta}_{mi} \Lambda \\ \Sigma &= \Lambda \Lambda^t \\ \tilde{\eta}_i &\sim MN(\mathbf{0}, \mathbf{I})\end{aligned}$$

with Λ a lower triangular matrix with real and non-negative diagonal values.

We implemented the estimation procedure based on the HMC method with the scale noncentered transformation using Stan (Stan Development Team, 2022). Based on our experience, we recommend to use an inverse Gamma distribution $\Gamma^{-1}(1, 1)$ for the variance σ^2 when the model includes only one random effect, and the inverse Wishart distribution $Inv\text{-}Wishart_M(\mathbf{I})$ distribution with $M > 1$ random effects for the variance-covariance matrix Σ (Gelman & Hill, 2006). We also suggest to use a weakly informative normal distribution (e.g., $N(0, 4)$) for the parameters $\theta_g, \beta_h, \gamma_{mn}$.

6. Model comparison

The application of the random-effects DyNAM requires the specification of the fixed and random effects included in the model. Statistics based on information criteria are crucial to support researchers in choosing which statistics have random effects. This complements the model specification process based on theory-driven arguments.

Drawing on the literature on the random effects model, we suggest using the Deviance Information Criteria (DIC), Widely Available Information Criteria (WAIC or Watanabe-Akaike Information Criteria), and Leave-One-Out Cross-Validation criteria with Pareto-Smoothed Importance Sampling (LOO-PSIS) (Vehtari *et al.*, 2017). The information criteria mentioned above aim to estimate the expected out-of-sample error using a bias-corrected adjustment of within-sample error (Gelman *et al.*, 2014). In the context of DyNAMs, the expected out-of-sample error is the prediction error for a relational event with a model that uses all events in the sequence except the one considered for prediction to estimate the parameters. A brute force approach would require re-estimating the model for each event in the sequence. Each information criterion uses a different bias-correction adjustment to get around the requirement of re-estimation using a model estimated with the complete event sequence as an input. Interested readers can refer to Vehtari *et al.* (2017) and Gelman *et al.* (2014) for a detailed description.

Two versions of the information criteria listed above have been computed using either the conditional or marginal version of the likelihood depending on the focus wished for the information criteria. The choice between the two depends on the specific application and whether the model should be predictive for events sent by actors presented in the data or events sent by new actors. The conditional likelihood predicts for senders presented in the data and thus incorporates the random coefficients in the analysis. The marginal likelihood predicts for new senders and thus integrates the random coefficients out as in, e.g., adaptive quadrature approximation using the MCMC samples proposed by Merkle *et al.* (2019). For random-effects DyNAMs, we suggest using the marginalized criteria because the comparison focuses on the variability explained by including random effects rather than the particular values taken by them.

7. Application to dynamics on an online community

Online social media platforms represent a rich source for the study of relational events (Lerner & Lomi, 2017; Mulder & Leenders, 2019). A large body of research has focused on the most prominent sites like Twitter, Facebook, Reddit, and Instagram to study the propagation of information (Bakshy *et al.*, 2015), content generation (Bhattacharya *et al.*, 2019), and political position (Hemphill *et al.*, 2013), among others topics. Social media niche sites represent a better scenario to study social phenomenon in-depth, beyond the analyses on actors centrality, as demonstrated in studies of virality (Hemsley & Kelly, 2019) and professional development (Marlow & Dabbish, 2014). In the latter, a social niche site is an online platform that caters to a specific audience (e.g., designers, academics, or knitters).

We investigated the dynamics of an online community for aspiring and professional digital and graphic designers called *Dribbble.com*.¹ The website was founded in 2009 and is a social network site to share and get feedback on user-made work (Scolere, 2019). Dribbble uses an invite-only membership system for new users and has a specific daily and monthly rate limit on the number of post images. As a consequence, the work posted on the platform has a high standard, and users develop a sense of belonging to an elite community (Wachs *et al.*, 2017). Users label the pieces of work with one or more tags describing the content and characteristics of the image.

Users can interact in different ways on the platform: they can follow other users, and like and comment on the work of other users. Time-stamped information on these interactions was collected using the platform API service.

After describing the data set, we illustrate how to apply the random-effects DyNAM to analyze the Dribble data. In particular, we outline the model specification and perform model selection using information criteria.

7.1 Description of the data set

We analyzed the sequence of like events that took place from 2015-03-10 to 2015-03-13 among the users. We chose a week in March (only the weekdays, Tuesday to Friday) to avoid the effects of seasonal and cyclic behavior observed in the data. The sequence of likes displayed a yearly seasonal cycle, with a peak in January and a decline in the last trimester of the year. The four days reflect the activity of a typical week during 2015, a year after the initial consolidation period of the platform (2009–2013), and before a second growth period that happened in 2016. We excluded from the data set the users belonging to a team as those might be employees of design companies who potentially interact with each other differently than nonaffiliated users.

We focused on published design pieces that use the tag “UI” (User Interface), the most frequent tag used throughout data collection. The subset allowed us to study the behavior of users working in similar design areas. We considered all the users who posted at least one new image, or were part of at least three like events or two follow events either as sender or receiver during the days considered for the analysis. We ended up with a data set with 1,080 users, 4,671 like events, 992 follow events, and 297 newly published design pieces.

For any point in time t , we defined the sender set ($\mathcal{S}(t)$) as the collection of users that created an account in the platform before time t and the receiver set ($\mathcal{R}(t)$) as the subset of users in $\mathcal{S}(t)$ that posted a shot before time t . The process state is defined as

$$\mathbf{y}(t) = (\mathbf{x}^{(l)}(t), \mathbf{x}^{(f)}(t), \mathcal{S}(t), \mathcal{R}(t))$$

where $\mathbf{x}^{(l)}(t)$ is the binary network of the past like events and, similarly, $\mathbf{x}^{(f)}(t)$ the binary network of past follow events between users.

7.2 Model specification

As an illustration, we use different specifications of the DyNAM choice submodel including both random and fixed effects. All the models have the same fixed-effects specification, and differ in the specification of the random effects.

The fixed-effect specification includes the following statistics:

$$\begin{aligned} s_1(i, j, \mathbf{y}) &= x_{ij}^{(l)} \\ s_2(i, j, \mathbf{y}) &= x_{ji}^{(l)} \\ s_3(i, j, \mathbf{y}) &= x_{ij}^{(f)} \\ s_4(i, j, \mathbf{y}) &= x_{ji}^{(f)}. \end{aligned} \tag{6}$$

The statistic s_1 takes value 1 if user i liked a piece of work posted by user j in the past, 0 otherwise. This inertia effect plays the role of the “intercept” in the model and captures the elemental behavior of updating existing ties. The statistic s_2 captures the tendency to reciprocate previous likes of other users. Thus, the statistics s_1 and s_2 model the dependence of like events on previous like events. The statistic s_3 represents the tendency of actors to like the pieces of work posted by users followed on the platform. Similarly, the statistic s_4 expresses the tendency to reciprocate follow events by likes events. These effects model the dependence of like events on previous follow events, an expected behavior in online social platforms beyond the essential endogenous process generated by the likes events.

Table 3. Model specifications used to analyze the like relational events. Model (1) is a standard DyNAM. Models (2) to (4) are DyNAMs with random effects.

Model	Formula
(1)	$f(i, j, \mathbf{y}, \boldsymbol{\beta}) = \sum_{h=1}^4 \beta_h s_h(i, j, \mathbf{y})$
(2)	$f(i, j, \mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\alpha}) = \sum_{h=2}^4 \beta_h s_h(i, j, \mathbf{y}) + \alpha_{1i} s_1(i, j, \mathbf{y})$ $\alpha_{1i} = \gamma_{10} + \eta_{1i}$ $\eta_{1i} \sim N(0, \sigma^2)$
(3)	$f(i, j, \mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\alpha}) = \sum_{h=2}^4 \beta_h s_h(i, j, \mathbf{y}) + \alpha_{1i} s_1(i, j, \mathbf{y})$ $\alpha_{1i} = \gamma_{10} + \gamma_{11} a_1(i, \mathbf{y}) + \eta_{1i}$ $\eta_{1i} \sim N(0, \sigma^2)$
(4)	$f(i, j, \mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\alpha}) = \sum_{h=2}^4 \beta_h s_h(i, j, \mathbf{y}) + \alpha_{1i} s_1(i, j, \mathbf{y})$ $\alpha_{1i} = \gamma_{10} + \gamma_{12} a_2(i, \mathbf{y}) + \eta_{1i}$ $\eta_{1i} \sim N(0, \sigma^2)$

We included random coefficients for the inertia effect as this is considered a baseline effect in the receiver choice model of DyNAMs. The simplest random-effects specification includes only the parameter γ_{10} . The resulting model (Model (2) in Table 3) assesses how much variation of the random inertia coefficients is present in the analyzed event sequence. To explain part of the variability of the random effects using actors’ characteristics, we subsequently added the following statistics:

$$a_1(i, \mathbf{y}) = \log\left(\sum_j x_{ij}^{(l)}(t_0) + 1\right)$$

$$a_2(i, \mathbf{y}) = t_0 - t_{i_c}. \tag{7}$$

The statistic a_1 is the out-degree activity of the actors at time t_0 , the initial time of the analyzed event sequence that corresponds to 2015-03-10. We used the logarithm of the out-degree because of the heavy-tailed and skewed distribution of the statistic. The log transformation reduces the effect of the highly active users on the parameter size. The statistic a_2 is the lifetime in years of a user’s account measured at time t_0 . The term t_{i_c} is the time at which the user i created the account in the platform, which predates t_0 , the beginning of the analyzed period, in each case.

To facilitate the interpretation and the numerical stability of the HMC, we standardized the statistics a_1 and a_2 by subtracting their mean and dividing by their standard error. This transformation helps to keep a general set of priors for the a_n coefficients where the parameters are scale-free. In other cases, the prior should consider the scale of the data in a sensitive way to avoid having a strongly informed prior that might heavily influence the posterior distribution. The inclusion of the statistics a_1 and a_2 one at the time led to Models (3) and (4) in Table 3.

Table 3 shows the considered model specifications for the utility of the receiver choice. Model (1) is the model without random effects and is the baseline model against which the random-effects DyNAM are compared. Model (2) adds random coefficients for the inertia effect to Model (1). Models (3) and (4) further add the statistics a_1 and a_2 , respectively, in α_{1i} to explain part of the variability of the random coefficients of the inertia statistic. It is worth mentioning that the substitution of α_{1i} in the utility formula $f(i, j, \mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\alpha})$ for Models (3) and (4) yields to a representation where the corresponding a_n statistics is part of an interaction with the s_1 statistic. This formulation is used in the implementation of the models.




Figure 1. Credibility intervals and posterior means for the parameters of the models defined in Table 3. The white dot with the thick border represents the posterior mean. The horizontal lines show the 50% (thick line) and 90% credibility intervals (thin line).

7.3 Results

The data were preprocessed using the R package `goldfish` (Stadtfeld & Hollway, 2022) that allows computing the change statistics for each event for the model formulation described in Section 7.2. We used an inverse Gamma $\Gamma^{-1}(1, 1)$ as a prior distribution for the random-effects variance σ^2 and place a $N(0, 4)$ prior on each β_n, γ_{1n} parameter for the estimation.

We checked the convergence of the HMC algorithm using diagnostic plots (e.g., trace plots) and statistics (e.g., rank normalized split- \hat{R} statistic, effective sample size) as described in Vehtari et al. (2021).

Figure 1 and Table 4 report the results obtained from 1,250 draws each from 8 chain of the posterior distribution.²

Figure 1 shows the central 50% and 90% credibility intervals and the means of the posterior distribution for the coefficients in Models (1) to (4). The white dot with the thick border represents the posterior mean. The horizontal lines show the 50% (thick) and 90% (thin) credibility intervals indicating the admissible range of coefficients values compatible with the data. The central 50% credible interval is defined by the first and third quantiles, while the central 90% credibility interval by the percentile 5% and 95% of the posterior distribution. The different colours indicate the estimates of the quantities described above for each of the Models (1) to (4).

Comparisons among the estimates suggest that the fixed effect coefficients, except the one for the inertia statistic, do not differ significantly across the four models. This result is expected since the model specifications presented in Table 3 introduce random effects and actors covariates to explain the variability of the inertia coefficient.

The inertia coefficient in Model (1) has a mean posterior of 1.73. In the actor-specific random effects Model (2), the average inertia coefficient has a mean posterior of 2.18, portraying an increase with respect to the fixed effect in Model (1). The average inertia value increases further to 2.49 when the model accounts for the out-degree activity of the users [Model (3)], but remains the same when the lifetime of an account is considered [Model (4)].

Table 4. Summary statistics of the posterior distribution, mean, standard deviation (SD) and credibility interval (CI) for the models defined in Table 3.

Coefficient		(1)	(2)	(3)	(4)
Inertia (β_1)	Mean	1.73			
	SD	0.04			
	90% CI	(1.67, 1.79)			
Average inertia (γ_{10})	Mean		2.18	2.49	2.18
	SD		0.07	0.07	0.07
	90% CI		(2.07, 2.29)	(2.38, 2.61)	(2.07, 2.29)
Recip. likes (β_2)	Mean	0.32	0.32	0.33	0.32
	SD	0.05	0.05	0.05	0.05
	90% CI	(0.23, 0.40)	(0.24, 0.41)	(0.24, 0.41)	(0.24, 0.41)
Tie (β_3)	Mean	0.43	0.47	0.47	0.47
	SD	0.05	0.05	0.06	0.05
	90% CI	(0.34, 0.52)	(0.38, 0.56)	(0.38, 0.56)	(0.38, 0.56)
Recip. follow (β_4)	Mean	0.89	0.76	0.75	0.76
	SD	0.04	0.04	0.04	0.04
	90% CI	(0.82, 0.96)	(0.69, 0.82)	(0.68, 0.82)	(0.69, 0.83)
log(outdeg) (γ_{11})	Mean			-0.60	
	SD			0.08	
	90% CI			(-0.72, -0.48)	
Time (γ_{12})	Mean				-0.08
	SD				0.07
	90% CI				(-0.19, 0.03)
σ^2	Mean		1.29	0.82	1.28
	SD		0.16	0.12	0.16
	90% CI		(1.04, 1.56)	(0.63, 1.03)	(1.04, 1.55)

Similar results are observed for the estimates of the posterior mean of the variance of the random effects σ^2 . The value of σ^2 in Model (2) is 1.29. The introduction of the log-outdegree [Model (3)] leads to a posterior mean of 0.82 and reduces the unexplained variance of the inertia effect by 36% ($= (1.29 - 0.82) / 1.29 \cdot 100$). No significant reduction is noted when controlling for the lifetime of an account in Model (4) ($\sigma^2 = 1.28$).

Figure 2 shows the posterior mean (white dot), the central 50% (dark grey), and 90% (light grey) credible intervals for each random coefficient η_{1i} describing the unexplained variance of the inertia effect for each user. In the plots, the users are displayed according to the increasing value of the posterior mean in Model (2).

Figure 2 indicates that the unexplained variability decreases from Model (2) to Model (3), as suggested by the smaller posterior means and ranges of the credibility intervals. No differences are observed when comparing Model (4) to Model (2). This result aligns with the observations above.

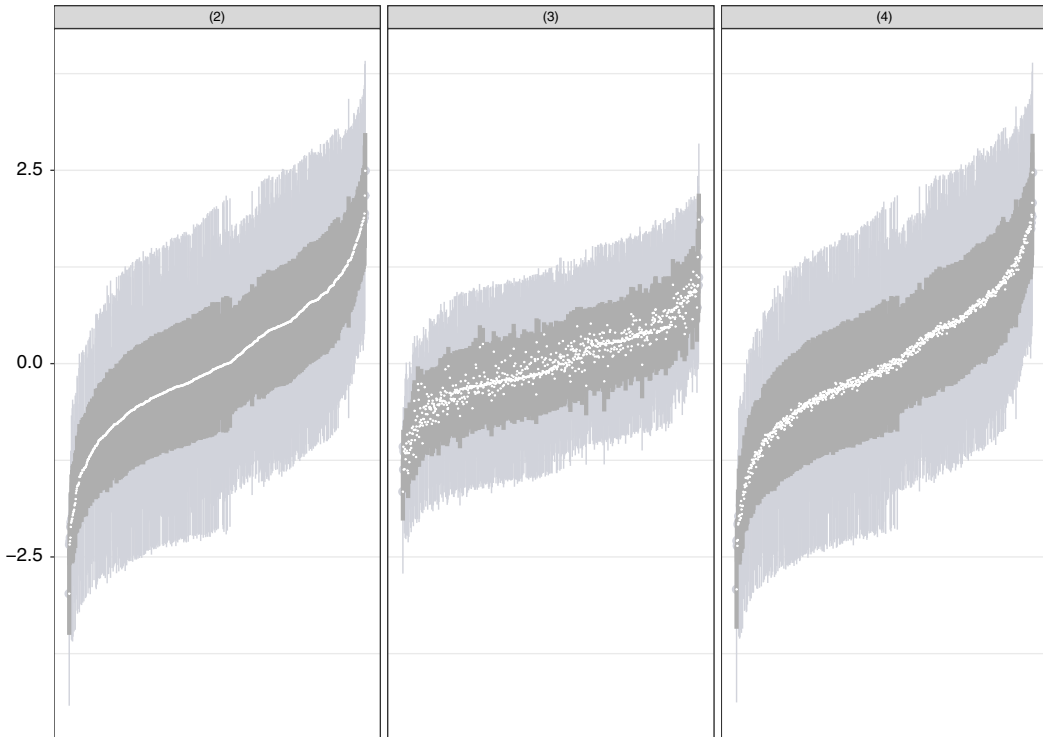


Figure 2. Credibility intervals for random effects η_{1i} by actors for the Models (2) to (4) described in Table 3.

We also note that there are differences in the length of the credibility intervals among the users. The analyzed data are unbalanced since users did not engage in the same number of events. For users with a few events, large credibility intervals are obtained, and the coefficient η_{1i} shrinks towards the global average tendency. In contrast, short credibility intervals and large deviations η_{1i} from zero are observed for actors engaged in a considerable number of events. This is a common finding in multilevel models (Gelman & Hill, 2006; Snijders & Bosker, 2011).

We conclude this section by formally comparing Models (1) to (4) based on the marginalized criteria introduced in Section 6 and illustrating the interpretations of the coefficients of the selected model. Table 5 shows the values of the information criteria for the four models. The models with the random effects have a better fit than the standard DyNAM [Model (1)], as suggested by the lower values of the criteria. In particular, the specification in Model (3) is the most adequate, as suggested by the inference of the posterior means and credibility intervals presented in the previous line.

Concerning the parameter interpretation of Model (3), the significant values of the fixed parameters in Table 4 can be interpreted as follows. The positive parameter for the reciprocity effect in the likes network (β_2) suggests a tendency to choose a receiver that liked the sender's artwork in the past. Similarly, the positive parameter for the tie effect in the following network (β_3) indicates that there is evidence that users tend to like pieces of work posted by users they follow on the platform. In addition, the positive parameter for the reciprocity effect in the following network (β_4) suggests a preference for receivers who followed the sender on the platform. For the random effects related to the inertia, the value of the average inertia coefficient (γ_{10}) is positive and significant, indicating a tendency to choose a receiver that the sender had liked artwork pieces in the past. Additionally, for the models with random effects, the negative value of the parameter related to the logarithm of the outdegree ($\gamma_{11} < 0$) indicates that low-active users are more likely

Table 5. Information Criteria (IC), expected log pointwise predictive density (elpd) and effective number of parameters (p) with their standard error (SE) for the Models in Table 3

IC	Measure	(1)	(2)	(3)	(4)
DIC	elpd	-30,612	-30,352	-30,327	-30,352
	SE	76	2,277	2,275	2,277
	IC	61,224	60,703	60,654	60,704
	SE	152	4,553	4,550	4,554
	p	4.03	4.92	5.96	5.97
	SE	0.02	0.29	0.32	0.31
LOO	elpd	-30,612	-30,353	-30,329	-30,354
	SE	76	2,277	2,275	2,277
	IC	61,224	60,707	60,658	60,708
	SE	152	4,554	4,550	4,554
	p	4.43	8.63	9.81	9.86
	SE	0.06	1.16	1.24	1.23
WAIC	elpd	-30,612	-30,353	-30,329	-30,354
	SE	76	2,277	2,275	2,277
	IC	61,224	60,707	60,658	60,708
	SE	152	4,554	4,550	4,554
	p	4.43	8.63	9.80	9.85
	SE	0.06	1.17	1.23	1.23

to send like events to the same receiver than more active users. Finally, the positive value of the inertia random effect variance σ^2 suggests users differ in their tendency to send like events to the same receiver. The interpretation of the parameters is conditional on a particular user initiating a like event at a certain point in time, as modeled by the rate submodel.

8. Conclusion

In this paper, we proposed an extension of the DyNAMs and developed a practical and appropriate implementation of parameter estimation for this model. Random-effects DyNAM draw on the literature on multilevel models and include random effects to control for actor or context unobserved heterogeneity in relational event data.

The model formulation includes monadic network statics or actors' covariates to explain part of the unobserved variability of a given network mechanism. Consequently, this new formulation has the advantage of providing a more precise estimation of the individual variations from the mean population importance parameters, especially for actors that send a few events (Gelman & Hill, 2006). Furthermore, the coefficients related to the monadic actors' covariates allow making inferences regarding the strength and direction in which actors deviate from the average tendency of the network mechanism by every unit of change in the given covariate. This is an advantage compared to the related model of DuBois *et al.* (2013).

We have described the random-effects DyNAMs for the multinomial choice model that determines the receiver of the events. However, the framework is flexible. It can be generalized to include random effects in the model specification for the actors' activity rate in an analogous way. Similarly, the model can be extended to analyze multiple sequences of events or more complex relational data. The proposed estimation method can be applied when the random effects are only used in one of the submodels. In that case, the separability assumption holds. If the random effects appear in both submodels, the parameters must be estimated jointly to consider the natural correlation between the random effects from the same actor in the different submodels. The development of the joint estimation of the parameters is the subject of future work.

A Bayesian estimation procedure based on HMC algorithms has been developed and implemented to estimate the parameters of the models. Compared to the more traditional Gibbs and Metropolis-Hastings algorithms, the HMC methods are more efficient and allow the exploration of the space of the posterior distribution more quickly. The current implementation applied to the empirical example, which represents a middle-size data set, took 17 hours to generate the draws of the posterior distribution for Model (2) and 9 hours for Model (3), where the model includes a monadic covariate to explain the variability. Models were run on a quad-core Intel Xeon E3-1585Lv5 processor (3.0–3.7 GHz). However, an additional limitation refers to the computation time scalability for large datasets, which represents an issue for the application of the model extension. A possible solution is to exploit the sparsity of the event network. The sparsity induces incremental updates of the statistics to a few actors. Our model extension could use additional strategies proposed to reduce the computation time considering the network event's sparsity (Vu et al., 2011b; Perry & Wolfe, 2013). A second alternative is using nested case-control sampling to reduce the data set size as proposed by Vu et al. (2017) and Lerner and Lomi (2020) for REM.

A natural progression of this work is extending the modeling framework to scenarios where the events sequence generative process display seasonal or cyclic behaviors. Some strategies have been developed for REMs that require prior knowledge either of the regular intervals that produce the seasonality (Amati et al., 2019) or the time window size where variations happen (Mulder & Leenders, 2019). Those strategies can also be implemented using the random-effects DyNAM introduced here, allowing empirical studies in scenarios where the data presents seasonal or cyclic behaviors without restricting the analysis to a subset of the data. Formulation of a DyNAM directly modeling temporal trends is the object of ongoing research.

Applications of the proposed extension required that researchers specify which statistics should be included in the model using random effects. Snijders and Bosker (2011, sec. 6.4) discussed this issue for multilevel models offering a digested compendium of possible problems and guidelines to follow during the model specification. Snijders and Bosker described the role of the substantive theory and the statistical considerations in the model specification. Regarding the substantive theory, it is advisable to include random effects for statistics that relate to hypotheses on unobserved heterogeneity or are plausible from subject-matter knowledge. For the statistical considerations, in the absence of prior knowledge, researchers in the absence of prior knowledge, researchers may follow a data-driven approach to select the network statistics to have random effects. In general, it is not advisable to include all statistics as random effects because the model would induce long computation times and might fail to converge. Limiting the number of random effects to a few may be necessary so that the model describes the event sequence satisfactorily but without unnecessary complications.

We illustrated the application of the random-effects DyNAM by analyzing relational events among the users of an online community of aspiring and professional digital and graphic designers. Based on the argument that users might behave differently on the platform, we focused on the inertia effect describing the repetition of events between the same sender and receiver. This is a baseline effect in the receiver choice model of DyNAMs. The analysis indicates that there are indeed differences between users. Part of the inertia variability can be explained by the users'

activity, with a lower likelihood of active users repeating events than less active users, but not by the lifetime of the users' account.

We used a subset of the dataset where known cyclic and seasonal patterns identified after an exploratory data analysis are considered. Consequently, the model results illustrate the model's applicability rather than a formal empirical study. Hence, the findings reflect the dynamics of the likes events during the analyzed week for users working in a similar design area. It is possible that popular tags may have similar event dynamics to the ones found here, considering that the dataset used the most frequent tag. The behavior might differ for tags used for niche communities of designers working in specified trends of art or types of designs.

The application shows that the model could be of particular value in social media contexts, where often large data sets can be collected while little is known about user characteristics. However, we believe that the new random-effects DyNAM could be fruitfully applied in a variety of relational event studies where behavioral heterogeneity between individuals can be expected.

Data availability. The data we used for the illustration are not publicly available. They were collected by Prof. Johannes Wachs (Vienna University of Economics and Business) who was kind enough to share them with us.

Acknowledgements. We thank the anonymous reviewers and the editor Prof. Tom A.B. Snijders for their careful reading and constructive comments. We also like to thank the members of the Social Networks Lab at ETH Zürich, attendees to Duisterbelt, and Prof. Marijtte van Duijn (University of Groningen) for feedback on the model definition and specification. We are also grateful to Prof. Johannes Wachs (Vienna University of Economics and Business) for sharing the data for the illustration of the model.

Funding statement. This research has received funding from The Swiss National Science Foundation (grants no. 10001A_169965, 10DL17_183008, and P2EZP1_188022).

Competing interests. None.

Notes

1 <https://dribbble.com/>

2 The R code used to estimate and compare the models presented in this section is provided as supplementary material.

References

- Amati, V., Lomi, A., & Mascia, D. (2019). Some days are better than others: Examining time-specific variation in the structuring of interorganizational relations. *Social Networks*, 57, 18–33.
- Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239), 1130–1132.
- Betancourt, M. (2017). A conceptual introduction to Hamiltonian Monte Carlo. arXiv preprint. [arXiv:1701.02434](https://arxiv.org/abs/1701.02434).
- Betancourt, M., & Girolami, M. (2015). Hamiltonian Monte Carlo for hierarchical models. *Current Trends in Bayesian Methodology with Applications*, 79(30), 2–4.
- Bhattacharya, P., Phan, T. Q., Bai, X., & Airolidi, E. M. (2019). A coevolution model of network structure and user behavior: The case of content generation in online social networks. *Information Systems Research*, 30(1), 117–132.
- Borgatti, S. P., & Halgin, D. S. (2011). On network theory. *Organization Science*, 22(5), 1168–1181.
- Box-Steffensmeier, J. M., Christenson, D. P., & Morgan, J. W. (2018). Modeling unobserved heterogeneity in social networks with the frailty exponential random graph model. *Political Analysis*, 26(1), 3–19.
- Butts, C. T. (2008). A relational event framework for social action. *Sociological Methodology*, 38(1), 155–200.
- DuBois, C., Butts, C. T., McFarland, D., & Smyth, P. (2013). Hierarchical models for relational event sequences. *Journal of Mathematical Psychology*, 57(6), 297–309.
- Gelman, A., & Hill, J. (2006). *Data analysis using regression and multilevel/hierarchical models*. Cambridge, UK: Cambridge University Press.
- Gelman, A., Hwang, J., & Vehtari, A. (2014). Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, 24(6), 997–1016.

- Hemphill, L., Otterbacher, J., & Shapiro, M. (2013). What's congress doing on twitter? In *Proceedings of the 2013 conference on Computer supported cooperative work - CSCW '13*. New York: ACM Press.
- Hemsley, J., & Kelly, R. D. (2019). Scratching a Niche: How smaller social media players such as dribbble reflect the viral phenomenon. *Social Media + Society*, 5(4), 205630511989005.
- Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn Sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1), 1593–1623.
- Lerner, J., & Lomi, A. (2017). The third man: Hierarchy formation in Wikipedia. *Applied Network Science*, 2(1), 24.
- Lerner, J., & Lomi, A. (2020). Reliability of relational event model estimates under sampling: How to fit a relational event model to 360 million dyadic events. *Network Science*, 8(1), 97–135.
- Luce, R. D., & Suppes, P. (1965). Preference, utility, and subjective utility. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology, III* (pp. 249–409). New York: Wiley.
- Marlow, J., & Dabbish, L. (2014). From rookie to all-star. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing - CSCW '14*. New York: ACM Press.
- McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in econometrics* (pp. 105–142). New York: Academic Press.
- McFadden, D., & Train, K. (2000). Mixed MNL models for discrete response. *Journal of Applied Econometrics*, 15(5), 447–470.
- Merkle, E. C., Furr, D., & Rabe-Hesketh, S. (2019). Bayesian comparison of latent variable models: Conditional versus marginal likelihoods. *Psychometrika*, 84(3), 802–829.
- Mulder, J., & Leenders, R. T. A. (2019). Modeling the evolution of interaction behavior in social networks: A dynamic relational event approach for real-time analysis. *Chaos, Solitons & Fractals*, 119, 73–85.
- Neal, R. M. (2011). MCMC using Hamiltonian dynamics. In S. Brooks, A. Gelman, G. Jones, & X.-L. Meng (Eds.), *Handbook of Markov chain Monte Carlo* (Vol. 2). Boca Raton, FL: CRC press.
- Papaspiliopoulos, O., Roberts, G. O., & Sköld, M. (2007). A general framework for the parametrization of hierarchical models. *Statistical Science*, 22(1), 59–73.
- Patison, K. P., Quintane, E., Swain, D. L., Robins, G., & Pattison, P. (2015). Time is of the essence: An application of a relational event model for animal social networks. *Behavioral Ecology and Sociobiology*, 69(5), 841–855.
- Perry, P. O., & Wolfe, P. J. (2013). Point process modelling for directed interaction networks. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75(5), 821–849.
- Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods* (Vol. 1). Thousand Oaks, CA: SAGE.
- Schweinberger, M. (2007). *Statistical methods for studying the evolution of networks and behavior*. PhD thesis, University of Groningen.
- Schweinberger, M. (2020). Statistical inference for continuous-time markov processes with block structure based on discrete-time network data. *Statistica Neerlandica*, 74(3), 342–362.
- Scolere, L. (2019). Brand yourself, design your future: Portfolio-building in the social media age. *New Media & Society*, 21(9), 1891–1909.
- Snijders, T. A. (1996). Stochastic actor-oriented models for network change. *The Journal of Mathematical Sociology*, 21(1-2), 149–172.
- Snijders, T. A. (2001). The statistical evaluation of social network dynamics. *Sociological Methodology*, 31(1), 361–395.
- Snijders, T. A., & Bosker, R. J. (2011). *Multilevel analysis: An introduction to basic and advanced multilevel modeling*. London, UK: SAGE.
- Stadtfeld, C. (2012). *Events in social networks: A stochastic actor-oriented framework for dynamic event processes in social networks*. Karlsruhe, Germany: KIT Scientific Publishing.
- Stadtfeld, C., & Block, P. (2017). Interactions, actors, and time: Dynamic network actor models for relational events. *Sociological Science*, 4(14), 318–352.
- Stadtfeld, C., & Hollway, J. (2022). *goldfish: Statistical network models for dynamic network data*. R package version 1.6.4, Social Networks Lab at ETH Zürich and the Graduate Institute Geneva.
- Stadtfeld, C., Hollway, J., & Block, P. (2017a). Dynamic network actor models: Investigating coordination ties through time. *Sociological Methodology*, 47a(1), 1–40.
- Stadtfeld, C., Hollway, J., & Block, P. (2017b). Rejoinder: Dynamics and the grounds for actor-oriented network event models. *Sociological Methodology*, 47b(1), 56–67.
- Stan Development Team (2022). Stan modeling language users guide and reference manual 2.29.2.
- Thiemichen, S., Friel, N., Caimo, A., & Kauermann, G. (2016). Bayesian exponential random graph models with nodal random effects. *Social Networks*, 46, 11–28.
- Tranmer, M., Marcum, C. S., Morton, F. B., Croft, D. P., & de Kort, S. R. (2015). Using the relational event model (REM) to investigate the temporal dynamics of animal social networks. *Animal Behaviour*, 101, 99–105.
- Van Duijn, M. A., Snijders, T. A., & Zijlstra, B. J. (2004). p2: A random effects model with covariates for directed graphs. *Statistica Neerlandica*, 58(2), 234–254.

- Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27(5), 1413–1432.
- Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., & Bürkner, P.-C. (2021). Rank-normalization, folding, and localization: An improved Rhat for assessing convergence of MCMC (with discussion). *Bayesian Analysis*, 16(2), 667–718.
- Vu, D., Lomi, A., Mascia, D., & Pallotti, F. (2017). Relational event models for longitudinal network data with an application to interhospital patient transfers. *Statistics in Medicine*, 36(14), 2265–2287.
- Vu, D. Q., Asuncion, A. U., Hunter, D., & Smyth, P. (2011a). Continuous-time regression models for longitudinal networks. In *Advances in neural information processing systems*, 24a.
- Vu, D. Q., Asuncion, A. U., Hunter, D. R., & Smyth, P. (2011b). Dynamic ego-centric models for citation networks. In *Proceedings of the 28th international conference on machine learning (ICML-11)* (pp. 857–864). New York: ACM Press.
- Wachs, J., Hannák, A., Vörös, A., & Daróczy, B. (2017). Why do men get more attention? Exploring factors behind success in an online design community. In *Eleventh international AAAI conference on web and social media*.
- Zappa, P., & Vu, D. Q. (2021). Markets as networks evolving step by step: Relational event models for the interbank market. *Physica A: Statistical Mechanics and Its Applications*, 565, 125557.