sociological science

Interactions, Actors, and Time: Dynamic Network Actor Models for Relational Events

Christoph Stadtfeld, Per Block

ETH Zürich

Abstract: Ample theoretical work on social networks is explicitly or implicitly concerned with the role of interpersonal interaction. However, empirical studies to date mostly focus on the analysis of stable relations. This article introduces Dynamic Network Actor Models (DyNAMs) for the study of directed, interpersonal interaction through time. The presented model addresses three important aspects of interpersonal interaction. First, interactions unfold in a larger social context and depend on complex structures in social systems. Second, interactions emanate from individuals and are based on personal preferences, restricted by the available interaction opportunities. Third, sequences of interactions develop dynamically, and the timing of interactions relative to one another contains useful information. We refer to these aspects as the network nature, the actor-oriented nature, and the dynamic nature of social interaction. A case study compares the DyNAM framework to the relational event model, a widely used statistical method for the study of social interaction data.

Keywords: social networks; DyNAM; relational event model; social interactions; network dynamics

Interpersonal interactions, such as meetings and conversations, are at the core of major social network theories. However, empirical research to date mostly focuses on the study of durable interpersonal relations, such as perceived friendships and trust or advice relations, rather than on processes of interpersonal interaction. In this article, we propose that the empirical study of interpersonal interactions can offer insights that go beyond the study of stable relationships and can complement our understanding of social network evolution and processes. To that end, we advance Dynamic Network Actor Models (DyNAMs) (Stadtfeld, Hollway, and Block 2017) to allow for the study of relational events of interpersonal interaction.

Theoretically, the relationship between social interactions and durable interpersonal relations is widely established. For example, Festinger, Schachter, and Back (1950) discussed how spatial proximity and thus opportunities for interaction affect the formation of stable social ties, whereas Heider (1958:189–90) argued that the distressing effect of interactions with friends in unbalanced situations may lead to the dissolution of friendship. Diverse theories about social influence (e.g., Friedkin 1998), about the structural advantage of network positions such as structural holes (Burt 2004) and structural embeddedness (Feld 1997), about homophily (Huston and Levinger 1978), and about social exchange (Emerson 1976) are explicitly or implicitly concerned with the role of interpersonal interaction.

From a theoretical perspective, interpersonal interactions can be understood as reflections of individuals' preferences, available interaction opportunities, and individuals' perception of the social network that they are embedded in. By observing patterns of interactions through time, researchers can thus gain deeper

Citation: Stadtfeld, Christoph, and Per Block. 2017. "Interactions, Actors, and Time: Dynamic Network Actor Models for Relational Events." Sociological Science 4: 318-352.

Received: March 10, 2017

Accepted: April 9, 2017

Published: May 15, 2017

Editor(s): Jesper Sørensen, Olav Sorenson

DOI: 10.15195/v4.a14

insights into the microlevel social mechanisms on the individual level that shape and reflect complex social systems (Hedström and Bearman 2009). We refer to this as the *actor-oriented nature* of interpersonal interaction.

Like stable social relationships, observations of interpersonal interactions are not independent but reflect underlying social structures, such as reciprocity, clustering, or preferential attachment. Studying the dependence between instances of interactions may thus open insights into some of the most fundamental social network mechanisms and, potentially, how those differ between individuals in different network positions. We refer to this feature as to the *network nature* of interpersonal interaction.

When studying social interactions through time, another important dimension of dependence unfolds that can reveal how sequences of interactions operate within short and longer time intervals. Reciprocal interactions, especially in communication networks, may typically occur in short time sequences, reflecting turn-taking in conversations. Preferential attachment or clustering, on the other hand, may operate with a longer "process memory" in which events in the distant past may be considered more prominently by actors. Studying interpersonal interaction may thus open new insights into the relative timing of social network mechanisms. We refer to this aspect as to the *dynamic nature* of interpersonal interaction. The DyNAM framework that we propose in this article is tailored to the dynamic nature, network nature, and actor-oriented nature of interpersonal interaction.

We believe that this is a promising period to move forward with the study of interpersonal interaction in social networks. The advent of electronic communication, social media, and human sensor technologies have brought about the possibility to collect a wealth of fine-grained social interaction data. Interdisciplinary researchers have called for using such data to enrich the empirical toolbox of social scientists (Lazer et al. 2009; Kitts 2014). Besides these newly emerging, technology-driven advances, for many years social networks researchers have been collecting archival network data about, for example, political and financial relations that often contain detailed information about the timing of relational interaction. Thus, in a variety of substantive disciplines, social network data are available as sequences of time-stamped traces of relational interaction to which we refer as *relational event data*, a term coined by Butts (2008).

In general, data on interpersonal interactions consist of at least three elements[1]: the event time, the event sender, and the event receiver. One can think of examples such as meetings ("Person A meets person B at time T"), favors ("Person A does person B a favor at time T"), phone calls ("Person A calls person B at time T"), or e-mails ("Person A writes an e-mail to person B at time T"). In this article, we will repeatedly return to the case of phone call sequences for illustrative purposes, and we analyze one such data set in an empirical case study towards the end of the article.

This article develops DyNAMs (Stadtfeld et al. 2017) further, an actor-oriented class of statistical models tailored to the analysis of time-stamped network data. While Stadtfeld et al. (2017) introduced DyNAMs for the study of undirected, durable coordination ties that are time-stamped, such as treaties between countries or romantic relationships, in this article we advance DyNAMs for the study of directed event sequences. The inferential task of these DyNAMs is, broadly, to explain who will send an event (event sender) to whom (event receiver) at what point in time (event time). The explanatory parameters in the model can relate to, for example, events in the recent or more distant past (A calls B, after B called A), ongoing relations between individuals (A calls B, because A and B are friends), or functions of individual characteristics (A calls B, because A and B have attributes Z in common). The model attempts to closely reflect the network nature, actor-oriented nature, and dynamic nature of interpersonal interaction.

The DyNAM represents the network nature of interpersonal interaction by explicitly modeling emerging network patterns over time as known from established network methods, such as the Exponential Random Graph Models (EGRMs) (Lusher, Koskinen, and Robins 2013) or Stochastic Actor-Oriented Models (SAOMs) (Snijders, van de Bunt, and Steglich 2010). Additional to the dependence on structures in the network of previous interaction sequences, the model allows for taking the larger, potentially complex social system into account. This can be represented by other, potentially coevolving networks of stable ties between individuals, common associations to organisations (represented in two-mode networks), or nodal, dyadic, and global covariates.

The DyNAM is in the conceptual tradition of SAOMs, in which network changes are modeled as choices of individuals (Snijders et al. 2010) reflecting the model's actor-oriented nature. Thus, it allows inference on individual preferences, the interaction opportunities, and, potentially, individuals' perception of the social system. In the course of this article, the actor-oriented nature of the DyNAM is theoretically and empirically compared to the most widely used, tie-oriented method for the analysis of time-stamped network data, the Relational Event Model (REM) (Butts 2008).

The dynamic nature of the model is reflected by the model not only concerning the order but also the absolute and relative timing of events. This is achieved by allowing the probability of events to vary by the amount of time that has passed since previous events through specific time windows. As a consequence, the statistical model needs to take periods of inactivity into account in a coherent way, as neglecting this can result in biased estimates. The importance of explicitly treating time with reference to short-term and long-term dynamics is discussed and will be illustrated in a detailed case study.

The remainder of this article is structured as follows. The DyNAMs for Relational Events section introduces Dynamic Network Actor models for relational events. In particular, it introduces the mathematical formulation that extends prior work by Stadtfeld and Geyer-Schulz (2011) and Stadtfeld (2012) to multivariate process states and processes with right-censored intervals. It further revisits the actor-oriented modeling paradigm and contrasts the approach to tie-oriented models such as the relational event model (Butts 2008). The Estimation section discusses estimation strategies for DyNAMs and REMs that are implemented in the R package `goldfish`. The Empirical Illustration of DyNAM and REM: Phone Calls in a Student House section presents a straightforward empirical case study in which sequences of phone calls between inhabitants of a student housing community are related to self-reported friendship ties. The DyNAM is further compared to the REM based on likelihood-based model fit, interpretation of parameters, and computational complexity. The Conclusions section discusses the future applicability and the potential impact of the DyNAM.

Dynamic Network Actor Models (DyNAMs) for Relational Events

The model presented here is in the tradition of stochastic actor-oriented models (SAOMs; Snijders 1996) applied to relational event data (Stadtfeld 2012). Originally, SAOMs have been developed to analyze the evolution of networks recorded as panel data. In recent years, however, SAOMs have been applied further as simulation models (e.g., Schaefer, Adams, and Haas 2013) and to cross-sectional data (Snijders and Steglich 2015). The model in this article extends and modifies SAOMs to take into account the qualitatively different nature of relational events in comparison to network ties.[2] In particular, we focus on relaxing assumptions about process myopia through time by allowing the inclusion of time-dependent statistics while taking into account the additional complications that come with this approach. In the following sections, we refer to this adaptation of SAOMs to event data as DyNAMs, a terminology first used in Stadtfeld et al. (2017).

The Explanandum

A relational event is defined as a triplet ω that includes the event time t_ω, the event sender i_ω, and the event receiver j_ω (in accordance with Butts 2008:159):

$$\omega = (t_\omega, i_\omega, j_\omega)$$

Observed relational event data consist of time-ordered sequences $\Omega = \{\omega_1, \omega_2, \dots\}$ of such triplets. The events in this series are typically not independent of one another, as they may contain the same individuals multiple times and relational events can, and often do, directly trigger one another (e.g., replying to or forwarding a received message).[3] However, the dependence between events is generally not a nuisance but is of direct interest to a network researcher, as this dependence can represent how people interact, react to others' actions, and are affected by the same exogenous factors.

The inferential task of the model is, then, to explain the timing, the sender, and the receiver of the relational event. Each of these can depend on previous events, other relationships between actors, actor attributes, global variables (like the time of the day) and much more. Jointly, all information that can influence which event transpires next (including information on past events) constitute the so-called "process state." We introduce this distinction between the event sequence Ω (the "dependent variable") and the process state (the "independent variable") for conceptual reasons. As will become apparent in the course of this article, this clear distinction simplifies modeling but without imposing any restrictions on the scope of the model. At the same time, it makes a lot of modeling assumptions directly tractable.

Working with Multivariate Process States

The process state entails all information about the relevant world from the perspective of the model at each time-point over the duration of the modeled period. In principle, it is of arbitrary complexity. For the case at hand, we formally define the process state $y(t)$ as a multivariate system at time t:

$$y(t) = \left(x^{(1)}(t), x^{(2)}(t), \dots, z^{(1)}(t), z^{(2)}(t), \dots, A^{(1)}(t), A^{(2)}(t), \dots\right), \quad (1)$$
$$y(t) \in \mathcal{Y}$$

The elements $x^{(\cdot)}$ represent networks or relational information of different kinds that can be expressed as matrices, the elements $z^{(\cdot)}$ represent nodal or global covariates that can be expressed as vectors or scalars, and the elements $A^{(\cdot)}$ represent node sets that correspond to the dimensions of the matrices and vectors.

The process state elements include all relevant information about the current process state and the relevant past (for example, all events sent from one individual to another up to time t, or all events sent within a specific time window before time t). It is worth emphasizing that by encoding all relevant information about the past in the process state via different process state elements, the model probability of observing an event only depends on the process state and nothing else. For this reason the model that we define can be understood as Markovian.

The process state can also include unmodeled, exogenously defined elements (e.g., advice relations between individuals or changing individual covariates). Dyadic variables and group affiliation will typically be expressed as networks, respectively. The elements $x^{(\cdot)}$ can thus be one-mode as well as two-mode networks. All networks, dyadic attributes, and nodal attributes are related to one or two node sets $A^{(\cdot)}$. If group affiliation of individuals, for example, is represented by a two-mode network $x^{(\cdot)}$, two node sets $A^{(\cdot)}$ will represent the two sets of individuals and groups. Nodal covariates $z^{(\cdot)}$ could then be related to those node sets and, for example, represent individuals' gender and the age of a group.

As the process state is a function of time, all elements can be either constant or subject to change. The process state can change because of exogenous and endogenous processes. Endogenous processes are the ones that are described by a specific model, and exogenous processes are those that relate to information updates without being modeled explicitly. The most basic endogenous update happens after each realized event, when the network $x^{(\cdot)}(t)$ that records all events up to time t is updated to include the event that just occurred. All parts of the process state that include information about the past event sequence are updated after an event transpired. An exogenous change of the process state might occur when, for example, the global variables are updated (e.g., seasonality or day time) or actors join or leave the network for unmodeled reasons. Process state elements that reflect certain aspects of the process history (e.g., indicate events that have happened within the last 24 hours) will further update the process state with a specified time-lag (e.g., 24 hours after an event has happened). Such time-lagged updates will be discussed in the Right-Censored Intervals section.

Figure 1 illustrates an exemplary process state. It consists of five nodes connected in two one-mode networks ($x^{(1)}$ and $x^{(2)}$) and one two-mode network ($x^{(3)}$) with

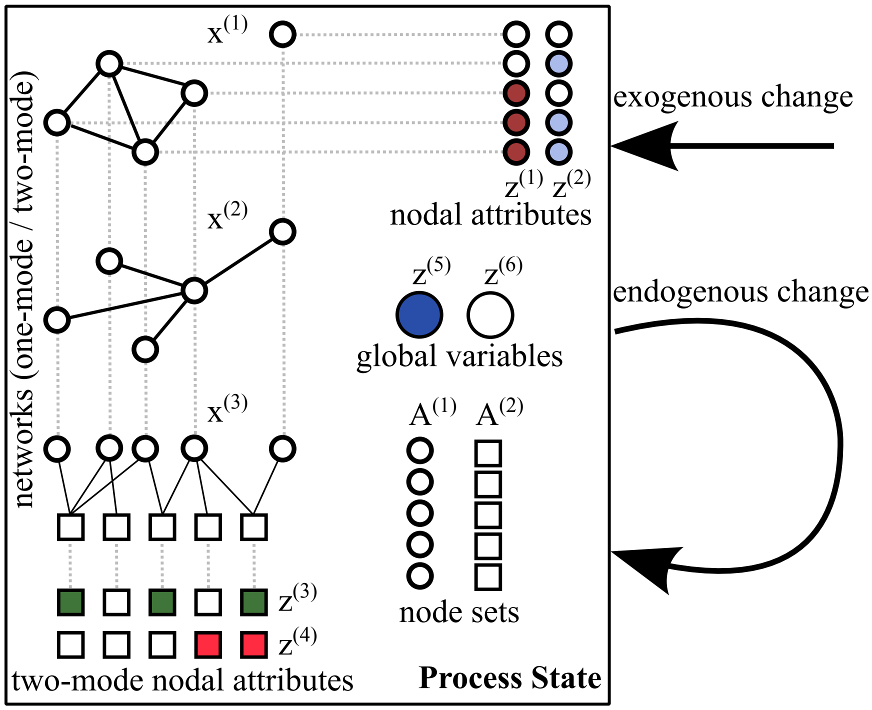

Figure 1: An exemplary multivariate process state.

five additional nodes on the second mode. Thus, two node sets are part of the process state ($A^{(1)}$ and $A^{(2)}$). One is indicated by circles, the other by squares. Further, the process state consists of two nodal attributes on each of the network modes and two global attributes ($z^{(1)}, \ldots, z^{(6)}$). Attributes are indicated by colors.

This conceptualization that includes all relevant model information in one process state simplifies working with increasingly complex network data, as practiced more and more in recent years and exemplified by recent studies that simultaneously analyze multiple networks (e.g., Boda and Néray 2015; Pál et al. 2016; Vörös and Snijders 2017) and, in particular, multimode or multilevel networks (Snijders et al. 2013; Wang et al. 2013; Lomi and Stadtfeld 2014; Hollway and Koskinen 2016).

An Actor-Oriented Model

Based on the process state y (for easier readability we omit the time indicator t when possible), we now define a probability model for the observed data Ω. In line with the basic SAOM (Snijders 1996), and similar to previous approaches (Stadtfeld and Geyer-Schulz 2011; Vu et al. 2011; Perry and Wolfe 2013), the proposed model is a two-step process with a linear predictor at its core. In the first step, the waiting time until an actor becomes active to send an event is modeled. In the second step, given that a specific actor becomes active, the recipient of this event is modeled. These two steps are assumed to be conditionally independent given the process state. Applied to an empirical example, intuitively the first step models who will make a phone call and at what time; the second step models who will be the receiver of that phone call.

Mathematically, the DyNAM models the waiting time between two subsequent events. The waiting time is modeled as a composite Poisson process in which each possible event observation $i \rightarrow j$ is associated with a Poisson rate $\psi_{ij}(y;\theta,\beta)$. The waiting times between subsequent events are thus exponentially distributed. Under the assumption that the two process steps (first, waiting time of a sender; second, discrete choice of a receiver) are conditionally independent given the process state y, we may express the first step as a Poisson process describing the activity of the actors and the second step as a probability distribution over all possible receivers:

$$\psi_{ij}(y;\theta,\beta) = \underbrace{\tau_i(y;\theta)}_{\text{select actor } i} \times \underbrace{p(i \rightarrow j; y, \beta)}_{\text{actor } i \text{ sends event to } j} \quad . \tag{2}$$

The waiting time τ_i of an actor i (step 1) depends on a parameter vector θ. The probability that an actor j is chosen as the event receiver ($p(i \rightarrow j)$; step 2) depends on a parameter vector β. Thus, the waiting time for each event (i.e., what will happen next) depends on the process state y and the two statistical parameters θ and β guiding the two modeling steps. These two steps are now introduced in detail.

As suggested by Snijders (2005:232), the actor rates τ_i can be modeled with an exponential link function that can be specified with arbitrary real statistic functions r_k evaluating the process state.

$$\tau_i(y;\theta) = \exp\left(\theta_0 + \sum_k \theta_k r_k(i,y)\right) \tag{3}$$

The parameter θ_0 refers to the general tendency of each actor to send events through time (i.e., the intercept). The effect functions $r_k(i,y)$ relate to differences between actors regarding their attributes or positions in the process state.

The individual activity rates τ_i model the timing until an actor is "getting active" as the sender of an event (e.g., picks up the phone to call someone), given the process state y. In Poisson models, the waiting time is exponentially distributed with parameters τ_i. In the language of event history modeling, $\tau_i(y;\theta)$ can be understood as the exponential hazard of actor i to get active—the hazard is independent of a waiting time δ.[4] An exponential waiting time model[5] is constructed by combining two functions:

$$\begin{aligned} f(\delta; y, \theta, i) &= \tau_i(y;\theta)e^{-\tau_i(y;\theta)\delta}, \\ S(\delta; y, \theta, i) &= e^{-\tau_i(y;\theta)\delta}. \end{aligned} \tag{4}$$

The first function f it the exponential probability density of an actor i to be active at time δ. The second function S is one minus the exponential cumulative distribution function and thus refers to the probability of actor i not being active within a period of length δ; in event history modeling, this function is called the survival function. The likelihood function (discussed in the Individual Activity Rates in the DyNAM section) will later be defined as a combination of density functions f (of actually observed events) and survival functions (of competing but nonrealized events).

In the models, multiple independent Poisson processes are competing in process state y, evaluating which actor in the actor set A will send an event next. The overall waiting time until any actor chooses to send an event is specified by an exponential distribution with parameter $\tau(y;\theta)$.

$$\tau(y;\theta) = \sum_{k \in A} \tau_i(y;\theta)$$

The probability of a specific actor i to be the first to be active in the current state y is proportional to the actor rates (hazards).

$$p(\text{actor } i \text{ is active next}; y, \theta) = \frac{\tau_i(y;\theta)}{\tau(y;\theta)}$$

Two of the three elements of the event triplet—event time and event sender—are hereby determined in the first modeling step.

In the second modeling step, the third element, the event receiver, is determined following a multinomial probability distribution (McFadden 1974). One possible interpretation is that receivers j are determined as a choice of the event sender i among all potential event receivers k given the process state y.

$$p(i \to j; y, \beta) = \frac{\exp\left(\beta^T s(i,j,y)\right)}{\sum_{k \in A \setminus \{i\}} \exp\left(\beta^T s(i,k,y)\right)} \tag{5}$$

The statistics $s(i,j,y)$ are functions of the process state that characterize the possible event $i \to j$. For example, they might indicate that j has previously sent an event to i (e.g., modeling the tendency to return a phone call) or that i and j are similar on some covariate (e.g., gender homophily). These statistics are in many ways comparable to statistics (or effects) used in classic SAOM or ERGM studies (Lusher et al. 2013). These statistics will be further explored in the following section.

The basic model that we describe here is almost identical to the formulation of Snijders' actor-oriented model (Snijders 2001)—a network change model was initially introduced for the analysis of actor-oriented processes in network panel data. However, some adjustments have been made to reflect the qualitatively different nature of interaction events and social ties, especially concerning multiple events on the same dyad and the explicit treatment of time, which allows us to relax of the assumption of strict myopia. Further discussion will be given in the Right-Censored Intervals section.

The Model Statistics

Equations (3) and (5) include statistic functions $r(i,y)$ and $s(i,j,y)$ that measure positions of nodes and ties in the process state. Thus, the selection of statistics in a model determines what a relational event process depends on and, consequently, which hypotheses can be tested in empirical analyses. The goal is to express that within certain structures the occurrence of an event will be more or less likely.

Micro-level mechanisms such as reciprocity, homophily, transitivity, and preferential attachment can be tested (for an overview, see Kadushin 2012 and Rivera, Soderstrom, and Uzzi 2010) by operationalizing them as a statistic functions.

Stadtfeld et al. (2017) provide an overview of possible mathematical specifications for the case of DyNAMs for undirected networks (coordination networks). The classification that they propose can be straightforwardly translated to directed networks and model specifications proposed for actor-oriented models (Snijders et al. 2010; Block 2015; Ripley et al. 2016) and ERGMs (Snijders et al. 2006; Lusher et al. 2013), and REMs (Butts 2008; Marcum and Butts 2015) can be applied well. Given the vast literature on statistics for actor-oriented models, we only conceptually introduce few model statistics that stand as examples for different effect classes.

- Effects that depend on past events:
  - The tendency to repeatedly send events to the same receiver within a time window. This effect can be operationalized as follows: $s_1(i, j, y) = x^{(1)}_{ij}$.
  - The tendency to reciprocate an event that has been received within a certain time window: $s(i, j, y) = x^{(1)}_{ji}$.

  Variable $x^{(1)}$ is a process state element of y that stores which events occurred within a certain time window as a weighted matrix. Matrix $x^{(1)}$ is updated through events as well as through lagged events (referring to updates at the end of a time window).

- Effects that depend on other networks:
  - The tendency to send events to a receiver that an actor is connected to in another network: $s_3(i, j, y) = x^{(2)}_{ij}$.
  - The tendency to send events to a receiver that has the same affiliations in a two-mode network: $s_4(i, j, y) = \sum_{h \in A^{(2)}} x^{(3)}_{ih} x^{(3)}_{jh}$.

  Variable $x^{(2)}$ and $x^{(3)}$ are process state elements of y that store information about another one-mode network (e.g., friendship) and a two-mode network (e.g., joint membership in organizations). $A^{(2)}$ is the set of nodes in the second network mode. Both networks are examples of process state elements that could be exogenously updated or be modeled by another event process.

- Effects that depend on actor attributes:
  - The tendency of actors to send more events if they score high on an individual attribute (a rate effect): $r_1(i, y) = z^{(1)}_i$.

  - The tendency to send events to receivers who score high on an individual attribute (a receiver choice effect):
    $s_5(i, j, y) = z_j^{(1)}$.

  Variable $z^{(1)}$ is a process state element of y that stores information about an ordered individual covariate (for example, individual attractiveness).

- Effects that depend on global attributes:

  - The tendency of actors to send more events if a global variable is high:
    $r_2(i, y) = z^{(2)}$.

  The global variable $z^{(2)}$ could, for example, be one at daytime and zero at night. A positive parameter would indicate that individuals are more active during the day.

- Effects that depend on a combination of the above:

  - The tendency to choose individuals as receivers who have been chosen recently by the sender and who the sender is indirectly connected to in another network:
    $s_6(i, j, y) = x_{ij}^{(1)} \sum_{h \in A^{(1)}} x_{ih}^{(2)} x_{hj}^{(2)}$.

  This effect could model whether friends of friends who have been called recently are more likely to be called again. Flexible combinations of the type of statistics above are possible.

In empirical analyses, these statistics can be calculated for each interval, combinations of the sender and potential receivers, and each effect. An important observation is that these model statistics cannot just be used for the estimation of DyNAMs but also for other types of network models such as the relational event model (Butts 2008). The main difference between the DyNAM and the REM is how the statistics are used to calculate probabilities, not the statistics themselves.[6] Practical considerations about model estimation are discussed in the Estimation section.

Right-Censored Intervals

In the introduction of the process state, it was considered that the process state may change because of exogenous processes and time-lagged updates. An example for the latter are cases in which nodes are allowed to show different activity patterns in the period after they have sent or received an event (i.e., explicitly modeling time heterogeneity of the process.) Within a certain time window, the process state would indicate that an actor is in a state of recent activity, allowing this to influence the subsequent probability to act. At the end of such an activity period, the process state is updated deterministically. Updates of the process state that do not follow endogenous events still constitute relevant model intervals—we refer to such intervals as *right-censored intervals*. Further examples of right-censored intervals are when exogenous information about unmodeled networks or nodal characteristics change, or when the set of actors is updated. In practical applications with many exogenous updates and detailed hypotheses about actor activity within time windows, the number of right-censored events can clearly outnumber the number of intervals that end with an endogenous event.

Although not necessarily obvious, it is important to consider those intervals in the estimation of models, especially for estimating time-dependent parameters. For example, assume we hypothesize that actors that were active within the past hour are more likely to send an event. In case we omit right-censored intervals in which no event was sent within this hour, we would tend to over-estimate the importance of this effect. This is because we would disregard information about all the cases in which actors did not act during this period—in a way, this would parallel an event history model in which cases were deleted when the dependent variable is 0 but a specific independent variable is 1. Thus, considering right-censored intervals is crucial for unbiased estimates of time-dependent parameters.

Right-censored and event intervals are illustrated in Figure 2 for a process of phone calls. The process state in this case consists of three nodes connected through a network representing all call events in the past (upper row), a network representing all calls within the past day (second row), and a nodal attribute that is subject to exogenous change (the color of nodes in the upper row). Nine time intervals are defined by the process (I to IX) within which the process state is constant. Change is brought about by endogenous updates that are due to phone calls occurring (intervals I, IV, and V), by exogenous change of the nodal covariates (intervals II and VII), and by time-lagged updates of the second network one day after a phone call was observed (intervals III, VII, and VIII). The final interval IX is right-censored by the end of the observation period. Six out of nine intervals are thus right-censored. In the estimation of parameters, these right-censored intervals need to be considered as they entail relevant information about model parameters.

Revisiting the Actor-Oriented Paradigm

Although the actor-oriented paradigm for the analysis of network evolution is, by now, firmly established, it is still instructive to review its foundational motivation, as discussed by Snijders (1996). Empirical tests of social theories often deduce associations between variables and test these in an established statistical framework, but this is not the most direct test of these theories. As noted *ibid.*:

> It would be preferable that the . . . deductions of the theory's implications be integrated with the statistical model that is used for the empirical test. Such an integration leads to a statistical model that is itself a direct expression of the sociological theory. . . . The integration of theoretical and statistical model is more complicated but it can lead to more stringent theory development because it requires a completely explicit theory and provides a much more direct test of the theory (Snijders 1996:149).

If we assume that network ties or, in our case, relational events emanate from actors that make conscious choices on whom to relate to, these decisions should be

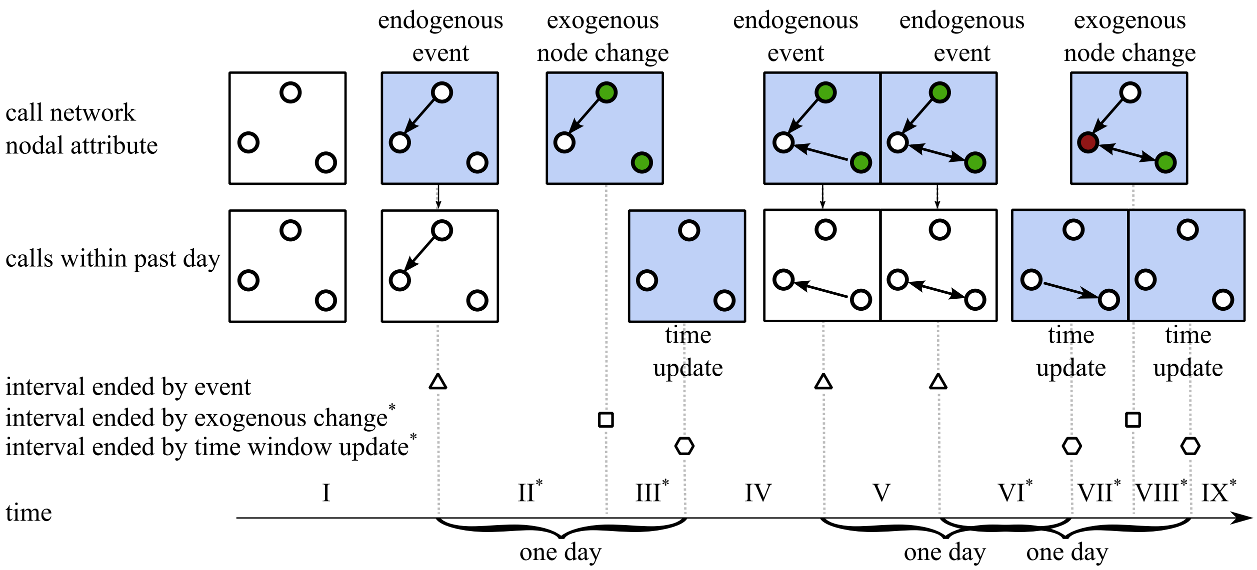

Figure 2: An exemplary process of call events that changes because of an endogenous process (call events change two network representations), a time window–related process (a network represents all phone calls that were made within the past day), and an exogenous process (update of a nodal attribute). The process state elements are shown in squares. Intervals that end because of exogenous changes or because of the update of time window representations (right-censored intervals) are marked with an asterisk.

explicitly modeled. This is the reason why the DyNAM, just like the SAOM, models what makes actors more or less likely to act and, once they act, whom they interact with. This explicit modeling of actions coming from individual actors originates in methodological individualism.

While the motivation underlying an actor-oriented approach is straightforward, its consequences are not obvious and require some further discussion that are most easily illustrated by comparison to the tie-oriented REM. The differences are that the REM contains one modeling step (which event will happen next), whereas the DyNAM contains two steps (who will send an event, and to whom; see Figure 3).[7] As a consequence, first, as opposed to the REM all modeling decisions are nested within actors (Block et al. 2016). This means that the probability to observe events in the REM only depends on its embedding in structures described by included effects. In the DyNAM, the probability to observe a specific event additionally depends on the other options an actor has to send events, as the different options for event receivers are compared by the sending actor. Thus, in an actor-based interpretation of the REM, a sender A decides whether it wants to send an event to a receiver B or not, whereas in the DyNAM, actor A decides whether it wants to send an event to actor B, rather than to any other actors C, D, E, et cetera.

Second, the interpretation of model parameters differs between the REM and the DyNAM. In the DyNAM, the time process and change process can be interpreted independently because those manifest in two separate modeling steps, whereas in the REM, these processes are combined in one modeling step. This means that in the DyNAM, we can interpret the parameters in the time choice function as who will be

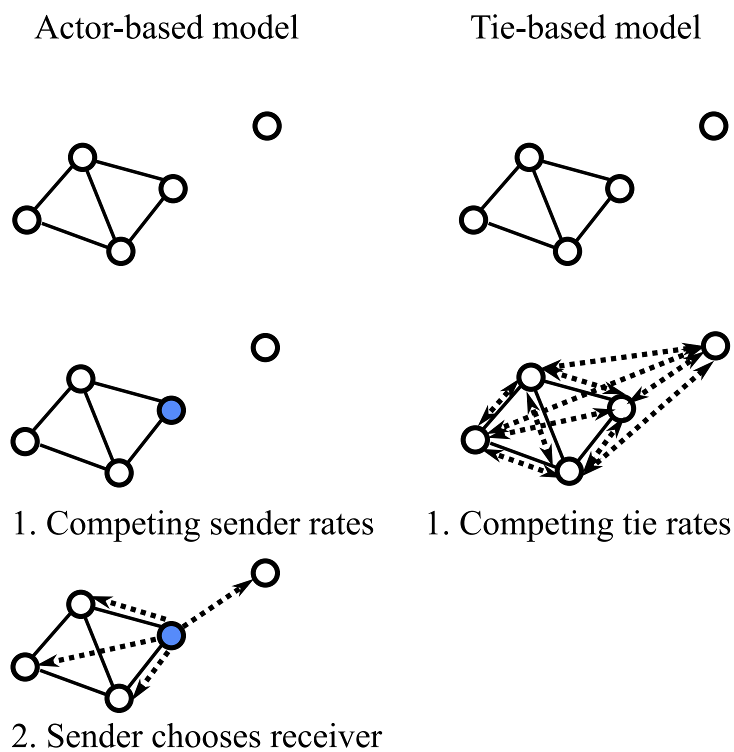

Figure 3: Comparison of actor-oriented and tie-oriented models.

active next and when and parameters in the receiver choice function as, "given that a certain actor is active," who is most likely to be chosen as a receiver. In the REM, both statistics that relate to whom will be active next and what will that person do must be interpreted net of the respective other model part. These differences are exemplified in the empirical example, which also discusses additional model evaluation criteria. Although these differences do not suggest that either model is "better," they show that the models differ inasmuch as the the nesting of decisions in actors and the interpretation of parameters.

Estimation

This section presents the estimation of parameters from data for the actor-oriented DyNAM and the tie-based REM. Compared to the literature, the latter is extended to incorporate right-censored intervals that stem from newly developed lagged window effects and exogenous process updates. All estimation routines have been implemented in the R package `goldfish` that is freely available for scientific use.[8]

The DyNAM consists of two subprocesses that are assumed to be conditionally independent, given the process state. The objective of the parameter estimation is to find the parameters $\hat{\theta}$ and $\hat{\beta}$ under which the observed data are as likely as possible. To this end, the likelihood of observed data needs to be defined. In order to maximize this, the calculation of the first and second derivative (or a useful approximation of it) of these likelihoods is necessary as well.

Individual Activity Rates in the DyNAM

This section introduces the estimation procedure of the first sub process that was defined in Equation (3), in particular, of the maximizing parameter vector $\hat{\theta}$. This first modeling step is concerned with the actor rates that determine two of the three elements of the event triplet: the time points t_ω and the event senders i_ω.

For each period observed one can now construct a likelihood function that makes use of the density function f (referring to actors who got active: observed events) and the survival functions S (actors who could have been active but were not: nonobserved events) in Equation (4). In right-censored intervals all potential events are non-observed. Let Ω be the set of all events, and Ω' the set of right-censored intervals. The union set is denoted by Ω^*. For each element ω_k in Ω^*, we define a time interval $\delta_{\omega_k} = t_{\omega_k} - t_{\omega_{k-1}}$. This is the length of an interval that is ended by an update of the process state. The set of actors that can act within a time interval is denoted by A_ω and the piecewise constant process state by y_ω.

The partial likelihood L^τ to observe the timing and senders of a specific sequence of events and right-censored intervals can then be expressed as:

$$
\begin{aligned}
L^\tau(\theta) &= \underbrace{\prod_{\omega\in\Omega}\left(\overbrace{\tau_{i_\omega}(y_\omega;\theta)e^{-\delta_\omega\tau_{i_\omega}(y_\omega;\theta)}}^{\text{observed event}}\overbrace{\prod_{k\in A_\omega\setminus\{i_\omega\}} e^{-\delta_\omega\tau_k(y_\omega;\theta)}}^{\text{nonobserved events}}\right)}_{\text{intervals ended by an event}}\underbrace{\prod_{\omega'\in\Omega'}\prod_{k\in A_{\omega'}} e^{-\delta_{\omega'}\tau_k(y_{\omega'};\theta)}}_{\text{right-censored intervals}}\\
&= \prod_{\omega\in\Omega}\left(\tau_{i_\omega}(y_\omega;\theta)\prod_{k\in A_\omega} e^{-\delta_\omega\tau_k(y_\omega;\theta)}\right)\prod_{\omega'\in\Omega'}\prod_{k\in A_{\omega'}} e^{-\delta_{\omega'}\tau_k(y_{\omega'};\theta)}\\
&= \prod_{\omega\in\Omega}\tau_{i_\omega}(y_\omega;\theta)e^{-\delta_\omega\sum_{k\in A_\omega}\tau_k(y_\omega;\theta)}\prod_{\omega'\in\Omega'} e^{-\delta_{\omega'}\sum_{k\in A_{\omega'}}\tau_k(y_{\omega'};\theta)}\\
&= \prod_{\omega\in\Omega}\tau_{i_\omega}(y_\omega;\theta)\prod_{\omega^*\in\Omega^*} e^{-\delta_{\omega^*}\sum_{k\in A_{\omega^*}}\tau_k(y_{\omega^*};\theta)}
\end{aligned}
$$

This can be understood as the likelihood of a Poisson process with exponentially distributed waiting times and right-censored intervals. The formula expresses that the likelihood of an observed time interval is for the left multiplier ("intervals ended by an event") the product of the exponential distribution density function (the hazard function) and the survival functions of all other rates (1 – distribution functions of the exponential distribution)—i.e., that actor i was getting active and not any other actor k. The right multiplier ("right-censored intervals") expresses that in the right-censored intervals no actor became active and it thus consists of a product of individual survival rates. As is customary in maximum likelihood estimation, its logarithm, but not the likelihood, is maximized:

$$
\begin{aligned}
\log L^\tau(\theta) &= \sum_{\omega\in\Omega}\log\left(\tau_{i_\omega}(y_\omega;\theta)\right) - \sum_{\omega^*\in\Omega\cup\Omega'}\delta_{\omega^*}\sum_{k\in A_{\omega^*}}\tau_k(y_{\omega^*};\theta)\\
&= \sum_{\omega\in\Omega}\left(\theta_0 + \sum_k \theta_k r_m(i_\omega, y_\omega)\right) - \sum_{\omega^*\in\Omega\cup\Omega'}\delta_{\omega^*}\sum_{k\in A_{\omega^*}}\tau_k(y_{\omega^*};\theta) \qquad (6)
\end{aligned}
$$

The maximum of the likelihood function is reached when the first derivative of the likelihood with regard to each of the statistics r_m equals zero. The first derivative of the likelihood is defined as:

$$\frac{\partial \log L^{\tau}(\theta)}{\partial \theta_m} = \sum_{\omega \in \Omega} r_m(i_\omega, y_\omega) - \sum_{\omega^* \in \Omega \cup \Omega'} \delta_{\omega^*} \sum_{k \in A_{\omega^*}} \tau_k(y_{\omega^*}) r_m(k, y_{\omega^*}) \tag{7}$$

For the partial derivative of $\log(L)$ with respect to the intercept θ_0, we define conveniently $r_0^{\tau}(k,y) = 1$ for any value of y and k. For this special case the first derivative reduces to

$$\frac{\partial \log L^{\tau}(\theta)}{\partial \theta_0} = |\Omega| - \underbrace{\sum_{\omega^* \in \Omega \cup \Omega'} \sum_{k \in A_{\omega^*}} \delta_{\omega^*} \tau_k(y_{\omega^*})}_{\text{expected number of events}} \tag{8}$$

where the first summand is the number of events observed and the second is the expected number of events over all periods (it is sum over the expectation of a sequence of Poisson distributions). The equation is zero and thus at maximum if the number of observed events equals the number of expected events over all event periods and right-censored periods. The interpretation of Equation (7) for the nonintercept parameters is somewhat similar, but there the observed and expected observations are weighted with the statistics functions. Equation (8) calls to attention that taking into account right-censored intervals is indeed important, as the right side is not invariant to the right-censored intervals in Ω'.
The second derivative is defined in Equation (9).

$$\frac{\partial \log L^{\tau}(\theta)}{\partial \theta_m \partial \theta_n} = - \sum_{\omega^* \in \Omega \cup \Omega'} \delta_{\omega^*} \sum_{k \in A_{\omega^*}} \tau_k(y_{\omega^*}) r_m(k, y_{\omega^*}) r_n(k, y_{\omega^*}) \tag{9}$$

The maximum likelihood estimate can then be found using, for example, an iterative Newton–Raphson procedure (Deuflhard 2004) based on the first and second derivative shown above. Alternative estimation procedures could be applied, such as a Fisher scoring method or other numerical procedures. The following section illustrates the Newton–Raphson logic in more detail.

Multinomial Receiver Choices in the DyNAM

The previous section explained partial likelihood inference related to the elements i_ω (sender) and t_ω (time) of the event triplet $\omega = (i_\omega, j_\omega, t_\omega)$. The second partial likelihood is related to the determination of the receiver j_ω and finding a likelihood maximizing parameter vector $\hat{\beta}$. Based on Equation (5), the likelihood of an event sequence Ω can be formulated as

$$L^{\beta}(\beta) = \prod_{\omega \in \Omega} \frac{\exp\left(\beta^T s(i,j,y)\right)}{\sum_{k \in A \setminus \{i\}} \exp\left(\beta^T s(i,k,y)\right)}$$

with variables $i = i_\omega$, $j = j_\omega$, $y = y(t_\omega)$, $A = A_\omega$. The event log likelihood for one event ω is given by

$$\log L^{\beta}_{\omega}(\beta) = \beta^T s(i,j,y) - \log\left(\sum_{k \in A\setminus\{i\}} \exp\left(\beta^T s(i,k,y)\right)\right). \tag{10}$$

We propose to estimate the parameter vector with a Gauss/Fisher scoring method as, for example, explained by Cramer (2003:36–38). A similar estimation routine that explicitly calculated the Hessian rather than approximating it with the negative information has been proposed by Stadtfeld (2012:45ff). The first derivative of a likelihood term with respect to parameter β_m (the m-th element of the score vector) is

$$\begin{aligned}\frac{\partial L^{\beta}_{\omega}(\beta)}{\partial \beta_m} &= p(i \to j; y, \beta)\left(s_m(i,j,y) - \sum_{k \in A\setminus\{i\}} p(i \to k; y, \beta) s_m(i,k,y)\right) \\ &= p(i \to j; y, \beta)\left(s_m(i,j,y) - \overline{s_m}(i)\right).\end{aligned} \tag{11}$$

Function $p(i \to j; y, \beta)$ was formulated in Equation (5). The first derivative of the log likelihood for one event ω is defined as

$$\frac{\partial \log L_\omega(\beta)}{\partial \beta_m} = s_m(i,j,y) - \overline{s_m}(i). \tag{12}$$

The abbreviation $\overline{s_m}(i)$ expresses the expected realized statistic m of an actor i's receiver choice given parameter β. For example, this could be the expected number of reciprocated relations that i is embedded in when choosing to send an event to one of the potential event receivers k in A with probability $p(i \to k; y, \beta)$ of Equation (5).

The Fisher information matrix F is the inverse of the expected matrix of second derivatives and can be utilized as an approximation of the latter in an iterative Newton–Raphson optimization. The advantage is that the Fisher information matrix is typically easier to calculate. For a detailed introduction we refer to Cramer (2003). For the receiver choice submodel of the DyNAM, an entry with indexes m and n of the Fisher information matrix is defined as follows:

$$F_{\omega;mn}(\beta) = \sum_{k \in A\setminus\{i\}} p(i \to k, y, \beta)\left(s_m(i,k,y) - \overline{s_m}(i)\right)\left(s_n(i,k,y) - \overline{s_n}(i)\right). \tag{13}$$

An updating step in the Fisher scoring method of a parameter β_t in iteration t uses the scores as well as the Fisher matrix to iteratively approximate the maximum likelihood estimate $\hat{\beta}$.

$$\begin{aligned} \beta^{t+1} &= \beta^t + F(\beta^t)^{-1} \frac{\partial \log L(\beta^t)}{\partial \beta^t}, \\ F(\beta^t) &= \sum_{\omega \in \Omega} F_\omega(\beta^t) \\ \log L(\beta^t) &= \sum_{\omega \in \Omega} L_\omega(\beta^t) \end{aligned} \tag{14}$$

Relational Event Model

The relational event model introduced by Butts (2008) is essentially a Poisson model that is equivalent to the one presented in the Individual Activity Rates in the DyNAM section. It is, however, somewhat more complex than the estimation of actor activity rates because the competing Poisson rates are defined for all sender–receiver combinations. In a network with n active nodes, n Poisson rates will be competing in the DyNAM but $n(n-1)$ Poisson rates in the REM. The rates for each sender–receiver combination $\{i, j\}$ can be defined based on a linear function (see Butts 2008:166) with a fixed intercept parameter γ_0 and additional parameters weighting structures in the process state.

$$\lambda_{ij}(y) = \exp\left(\gamma_0 + \sum_k \gamma_k s_k^\lambda(i, j, y)\right) \tag{15}$$

The log likelihood then looks similar to Equation (6). Functions s_k are model statistics like the ones introduced in The Model Statistics section. Their effect on the process dynamics is described by parameter vector γ that is subject to statistical inference. Other than in its original formulation in Butts (2008), we incorporate right-censored intervals throughout the process and thus allow exogenous as well as lagged process changes.

$$\log L^\lambda(\gamma) = \sum_{\omega \in \Omega} \left(\gamma_0 + \sum_k \gamma_k s_k^\lambda(y_\omega, i_\omega, j_\omega)\right) - \sum_{\omega^* \in \Omega \cup \Omega'} \delta_{\omega^*} \sum_{\{k,l\} \in A_{\omega^*}} \lambda_{kl}(y_{\omega^*}) \tag{16}$$

The implemented estimation routine is a Newton–Raphson optimization and resembles the one presented in the Individual Activity Rates in the DyNAM section. Only the computational and data complexity is higher, as more Poisson rates are competing. The log likelihood $\log L^\lambda(\gamma)$, however, is the complete process log likelihood whereas in the case of the actor-oriented models, the complete process log likelihood is determined as a sum of the two partial log likelihoods $\log L^\tau(\theta)$ and $\log L^\beta(\beta)$ (Eqs. [6] and [10]).

Partial Likelihood Inference and Semiparametric Models

One key advantage of parametric models such as the ones discussed above (DyNAM and REM) is that it is possible to use them for process simulations. Thereby, one can, for example, generate simulation-based goodness-of-fit criteria as know from ERGMs (Lusher et al. 2013) and SAOMs (Lospinoso 2012).

An alternative to parametric models such as the ones shown above are semiparametric models (achieved by integrating out the rates and focusing on the estimation of ordered models and a subset of non–time-dependent effects) (Butts 2008; Stadtfeld et al. 2017). This is reasonable for processes in which no right-censored intervals exist but is more problematic if this cannot be assumed.

Another alternative to full parametric models are Cox models (Cox and Oakes 1984). A theoretically plausible feature of Cox models is that distribution assumptions about the hazards, such as exponentially distributed waiting times in the case of DyNAM and REM, are not necessary. Rather, under the assumption of proportional hazards, Cox models are unconcerned about the actual underlying distribution. Perry and Wolfe (2013) proposed proportional hazard estimation for event data. This is reasonable as long as the model still assumes piecewise stable process states as in this article; this means that Cox models also need to consider right-censored intervals. Once effects are taken into account that, for example, exponentially weight the past,[9] the proportional hazards assumption is violated and it is unclear in which biases this would result.

Empirical Illustration of DyNAM and REM: Phone Calls in a Student House

The Social Evolution Data

As an empirical illustration, we analyze the "Social Evolution" data set of the Massachusetts Institute of Technology Media Lab, a rich data set of 84 individuals living in a U.S. student house. The data were originally collected to assess the "health state" of a social community, particularly by observing changes in physical health but also health-related norms (Madan et al. 2012). Here, we use the newly developed `goldfish` software to analyze a sequence of phone calls between the inhabitants (automated, time-stamped data collection) and relate those to change in a friendship network (based on a longitudinal survey at two points in time). This reflects our introductory comment about how interaction studies may complement network studies about stable relations, such as friendship.

Figure 4 shows the phone call network within the first two months of the data collection (black arrows, excluding isolates). Self-declared friendship relations are illustrated by blurry blue ties in the background. Node color relates to the grade type of individuals who are students at different stages of their respective study programs. On visual inspection it seems that friendship relations are related to phone call interactions and that the network exhibits a certain level of grade-type homophily. We make use of the DyNAM to test the effects of such descriptive findings in a dynamic framework. Not only do we present results of the DyNAM but we compare those to the relational event model (Butts 2008), in which right-censored intervals are additionally taken into account.

The process state that we define to model these data consists of six elements:

- a weighted network that measures who has called whom in the past ("callNetwork")

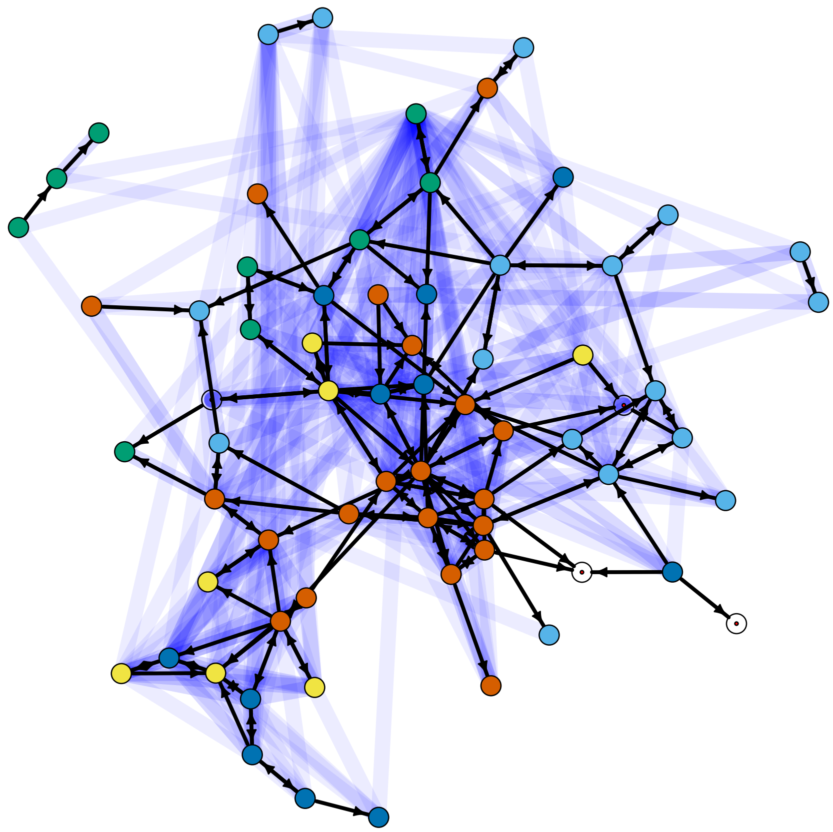

Figure 4: The call network (without isolates) of the first two months of data collection. Friendship relations are indicated by shadowed ties in the background. The node color relates to the grade type of individuals; missing values are indicated by nodes with a central dot.

- a weighted network that measures who has called whom in the past hour ("callNetworkPastHour")
- a network that expresses who has nominated whom as a friend ("friendship")
- an actor variable that expresses the floor on which individuals live in the student house ("floor")
- an actor variable that expresses the grade type of individuals ("gradeType")
- an actor set representing 84 individuals

The model statistics refer to these process state elements. The first three elements change over time whereas the last three are constant.[10] The second element ("callNetworkPastHour") is changed through events (ties are increased by one) as well as through lagged updates (values are decreased one hour after a phone call took place). The short names in parentheses above will be used in the following model specification as well as in the Results section.

Empirical Model Specifications

The specification of the empirical case study is kept straightforward and consists of effects that relate to the general activity of actors (the DyNAM rates) and the propensity to choose certain receivers (the DyNAM choice model). In the REM, all parameters simultaneously affect the timing and sequence of events, without a separation into two submodels. Mathematical details about the model specifications are provided in Appendix A of the online supplement. The naming conventions are in line with the short names of effects in the SIENA software (Ripley et al. 2016) for the estimation of actor-oriented models for network panel data.

Five effects are concerned with the general activity of actors in both models. In case of the DyNAM, these five effects specify the sender rate shown in Equation (3). The most basic one is an intercept that models

1. the general activity of actors to make phone calls ("intercept")

We further test how the position of nodes in the process states relates to their activity, in particular whether the propensity to make a phone call it is affected by the following:

2. having made phone calls within the past hour ("egoX recentCallsSent")
3. having received phone calls within the past hour ("egoX recentCallsReceived")
4. having more friends ("outdegreeX friendship")
5. having called more people in the past ("outdegreeX callNetwork")

Eleven effects are concerned with the choice of a call receiver. These effects constitute the specification of the receiver choice submodel in Equation (5). They test the following:

6. whether individuals tend to call those they have called before ("outdegree callNetwork")
7. whether individuals tend to call those they have called in the past hour ("outdegree callNetworkPastHour")
8. whether individuals tend to call those who have called them before ("reciprocity callNetwork")
9. whether individuals tend to call those who have called them in the past hour ("reciprocity callNetworkPastHour")
10. whether individuals tend to call those who have been called by many others ("inPop callNetwork")
11. whether individuals tend to call those who have been nominated as a friend by many others ("inPop friendship")
12. whether individuals tend to call those who have been called by past receivers of their own past phone calls ("transitivity callNetwork")

13. whether individuals tend to call friends of their friends ("transitivity friendship")

14. whether individuals tend to call others who live on the same floor ("sameX floor")

15. whether individuals tend to call others who are in the same grade type ("sameX gradeType")

16. whether individuals tend to call their friends ("X friendship").

Effects 2, 3, 7, and 9 are related to changes within a one-hour time window. Effects of that kind (e.g., the tendency of actors to make phone calls if they have called someone recently) were introduced as time window effects in Stadtfeld et al. (2017).

Results

Results of the case study are presented in Table 1. The leftmost column includes estimates of the DyNAM rate model, the center column those of the DyNAM receiver choice model, and the right column the estimates of a REM that was specified in the same way. Below the table, the log likelihoods are shown as well as the CPU time that was needed for the estimation. In case of the DyNAM, the joint log likelihood of both sub models is printed below the center column. The "interval" log likelihood relates to the full models as described in this article, whereas the "sequence" log likelihood relates to a simpler model in which only the order of 1,092 phone call events was modeled without taking into account timing and right-censored intervals. The Biases When Ignoring Right-Censored Intervals discusses biases in such sequence models in more detail.

In the DyNAM, both submodels can be interpreted independently because it assumes conditional independence of the two submodels given the process state.

The five DyNAM rate effects (effects 1–5) can be interpreted as "what affects the tendency of individuals to make phone calls," irrespective of who they are calling. The intercept effect models the general tendency of individuals to make phone calls. The intercept can be loosely interpreted as follows: Individuals who have not made and received calls recently, have no friends, and have in general never made a phone call (the model statistics of parameters 1–4 are thus zero) have an expected waiting time until they make a phone call of 10.7 days.[11] The other four rate parameters are positive and can thus be interpreted as all increasing the tendency of individuals to make phone calls (and thus decreasing the expected waiting time between subsequent calls). For example, our model results suggest that individuals who have made and received five phone calls within the past hour, have five friends, and have called five people ever before will only have an expected waiting time until the next phone call of 13 minutes.[12] DyNAM rate effects sizes can be interpreted in terms of relative probabilities. For example, one could loosely state that with each additional person someone has called within the past hour, the expected waiting time decreases by about 45 percent, all else being equal.[13] At the end of the one-hour window following a phone call a right-censored interval concludes the higher activity period of an actor.

Table 1: Results of the case study using DyNAM and REM.

		DyNAM		REM
		(1)	(2)	(3)
#	Effects	$\hat{\theta}$ (s.e.)	$\hat{\beta}$ (s.e.)	$\hat{\gamma}$ (s.e.)
1	intercept	−13.74†		−14.78†
		0.05		0.14
2	egoX recentCallsSent	0.59†		0.55†
		0.04		0.03
3	egoX recentCallsReceived	0.53†		−0.27†
		0.05		0.05
4	outdegreeX friendship	0.03†		−0.04†
		0.00		0.01
5	outdegreeX callNetwork	0.26†		−0.03
		0.01		0.02
6	outdegree callNetwork		−4.09†	−4.97†
			0.16	0.13
7	outdegree callNetworkPastHour		−2.06†	0.79†
			0.19	0.11
8	reciprocity callNetwork		0.24	0.38†
			0.14	0.11
9	reciprocity callNetworkPastHour		4.01†	4.67†
			0.32	0.13
10	inPop callNetwork		−0.10∗	0.01
			0.05	0.03
11	inPop friendship		−0.17†	−0.06†
			0.02	0.01
12	transitivity callNetwork		0.30∗	−0.16∗
			0.12	0.07
13	transitivity friendship		0.02	0.14†
			0.04	0.02
14	sameX floor		−0.24∗	−0.98†
			0.12	0.08
15	sameX gradeType		0.13	−0.17∗
			0.12	0.08
16	X friendship		2.07†	1.52†
			0.18	0.12
	Log likelihood (submodel)	−14,165.94	−1,319.38	−14,800.81
	Log likelihood (interval)		−15,485.32	−14,800.81
	Log likelihood (sequence)		−4,987.42	−4,513.31
	CPU time (seconds)		18.11	2,111.69

Column (1) includes the DyNAM sender rate, column (2) the DyNAM receiver choice, and column (3) the REM tie rate. Standard errors are reported below the estimates. The data set includes 1,092 phone call events and 1,093 right-censored intervals between 84 actors. The likelihood ratio per interval is 1.368 (REM over DyNAM, considering both events and right-censored intervals).
∗ $p < 0.05$, † $p < 0.01$

The eleven receiver choice effects (effects 6–16) can be interpreted as "what affects the choice of who will be called". Given that a certain individual makes a phone call at a certain point in time (modeled by the rate), the results show that it is very likely that a receiver will be chosen that has been called before (the negative "outdegree callNetwork" effect 6 suggests that individuals keep their outdegree low), particularly if the receiver has already been called within the past hour (negative "outdegree callNetworkPastHour" effect 7). Further, we find no evidence that individuals tend to call those who have called them before at any point in the past ("reciprocity callNetwork" effect 8); however, there is a strong effect if they have been called by that person within the past hour (positive "reciprocity callNetworkPastHour" effect 9). There is strong evidence that people tend to call their friends ("X friendship" effect 16). The "transitivity callNetwork" effect is positive, which means that individuals tend to call those who they are connected to through one or several two-paths in the phone call network. There is, however, no evidence that being connected through a friendship two-path increases the probability of choosing a receiver ("transitivity friendship") and thus calling friends of friends who are not friends (we control for the direct friendship effect already). The two effects relating to indegree popularity in the call network and in the friendship network (10:"inPop callNetwork"; 11:"inPop friendship") are negative, which indicates that the process is characterized by tendencies towards a balanced indegree distribution rather than a strong centralization. Only one of the two nodal effects is significant, and we find evidence for calling people who live on a different floor (14: "sameX floor"). An interpretation as relative probabilities is again possible. For example, our findings suggest that individuals who are friends are 7.9 times more likely to call each other, all else being equal.[14] It is important to note that the interpretation of effects within the receiver choice model are all net of each other. For example, we descriptively see that on an aggregated level, people have a higher tendency to call others who live on the same floor. Once we control, however, for the other parameters such as those that model that friendship networks are stronger within the floor network, we find a negative residual effect for floor homophily.

The interpretation of REM parameters is somewhat more involved as all parameters affect the waiting times between subsequent events as well as the tendency that an event of a specific sender–receiver pair is occuring next. The REM is tie-oriented, which means that all tie Poisson rates are competing at all times. In an actor-oriented model, the receiver choice effects are nested within actors and alternatives are thus only compared for the actor that actually decides to make a phone call. Some estimates of the REM indeed have a different sign as compared to the DyNAM, which illustrates the differences in interpretation. For example, the "outdegreeX friendship" effect is negative. This could relate to the significant "transitivity friendship" effect in the model: individuals with many friends tend to be embedded in more transitive friendship clusters within which phone calls are more likely and more frequent. The individual "outdegreeX friendship" effect thus potentially balances out the structural "transitivity friendship" effect. In the DyNAM, both effects can be interpreted independently, as they are part of different submodels. We could thus straightforwardly conclude that individuals with many friends are more likely to make phone calls—an observation that is supported by descriptive analyses of the data set but is difficult to infer from the REM results. Similar to the DyNAM, we can also interpret parameters in terms of relative probabilities; however, now we need to take the tie-oriented nature into account. The probability that the next event in the sequence will be a phone call between any pair of friends is 4.6 times more likely than a phone call between any pair of individuals who are not friends.[15]

Overall, it can be said that the REM has a significantly better fit regarding log likelihood (higher values are better). The likelihood ratio per event is 1.37, which means that the REM is on average 37 percent more precise for any event. This difference is explained by the fact that in the REM, all effects model waiting times and the event sequence. Further, all tie rates are defined as competing waiting time processes and their evaluation is not nested within actors. The two submodels of the DyNAM are either concerned with actor waiting times or with with the choices of event receivers. A drawback of the higher precision of the REM is the higher computational complexity, which led to a CPU time that is a factor 60 higher. The following two sections explore the model likelihoods and the computational complexity in more detail.

Comparison of Event Likelihoods

By zooming in on the likelihood of single events, researchers can gain deeper insights into where the DyNAM performs well and where there are differences between the actor-oriented DyNAM and the tie-oriented REM. Figure 5 shows the likelihood ratios for each of 1,092 events in two submodels that ignore right-censored intervals and are just concerned with ordered event sequences (like the earlier models proposed by Butts 2008 and Stadtfeld 2012). As some of the likelihood ratios (in particular for the case of unlikely events) have extremely high values of more than a thousand, the likelihood ratios were logged below and above the baseline. It can be seen that the REM performs better in most events. The average likelihood ratio of non-censored intervals (intervals ending with a phone call event) is 1.95 per event in favor of the REM, whereas the likelihood ratio of right-censored intervals is 1.09 in favor of the DyNAM. A closer look at two outliers better explains under which circumstances each model performs better.

Outlier 1 in Figure 5 is an event that is considered rather unlikely by both models (only 8 times more likely than choosing a random tie in the REM and quite impossible in the DyNAM, with just 0.07 of the random choice likelihood). The event likelihood of the REM, however, is more than 1,000 times higher than the DyNAM event likelihood. The event describes a phone call at 3:30 a.m. on a Monday morning between two individuals who have been calling each other regularly (but not within the last hour). The DyNAM probability of the call sender to make a phone call is extremely small because five other individuals have been calling each other within the past hour. The most likely event from a DyNAM's point of view would be that one of those five makes another phone call, given that "egoX recentCallsSent" and "egoX recentCallsReceived" are the dominant effects in the sparsely specified DyNAM rates. The REM, however, takes all effects into account when modeling waiting times between events. It can therefore factor into

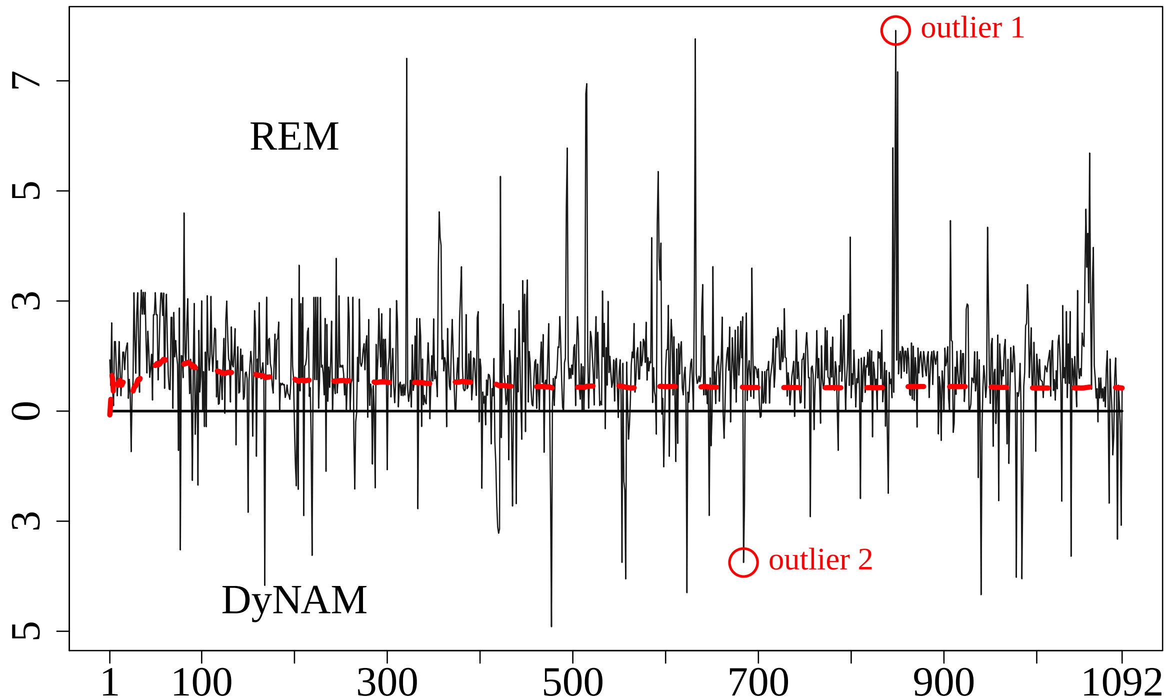

Figure 5: Likelihood ratios (logged for better readability) per event in an ordered model that ignores right-censored intervals and just models sequences. Events above the 0 line favor the REM, events below favor the DyNAM. Two outliers are exemplarily discussed in the text. The dashed line shows for each event the average likelihood ratios of all prior events.

the waiting time model that the two individuals are friends and have called each other before. Those two parameters are only considered in the DyNAM conditional on the probability that the actual sender makes a phone call, which is extremely unlikely given recent activity.

Outlier 2 in Figure 5 is an event that is considered rather likely in both models (149 times more likely than choosing a random tie in the DyNAM). The DyNAM event likelihood is 42 times higher than the REM event likelihood. Here, the DyNAM rates put a lot of weight on the actual call sender, as this person has a very high outdegree in the friendship network and in the call network. Both effects are associated with a high activity of actors. In case of the REM, these effects negatively contribute to the activity of actors net of all other effects (e.g., the "X friendship" effect).

The comparison of event likelihoods can help researchers improve the fit of their empirical models. Our exemplary analyses, for example, showed that the specific, exemplary DyNAM performs poorly in case the three rate parameters do

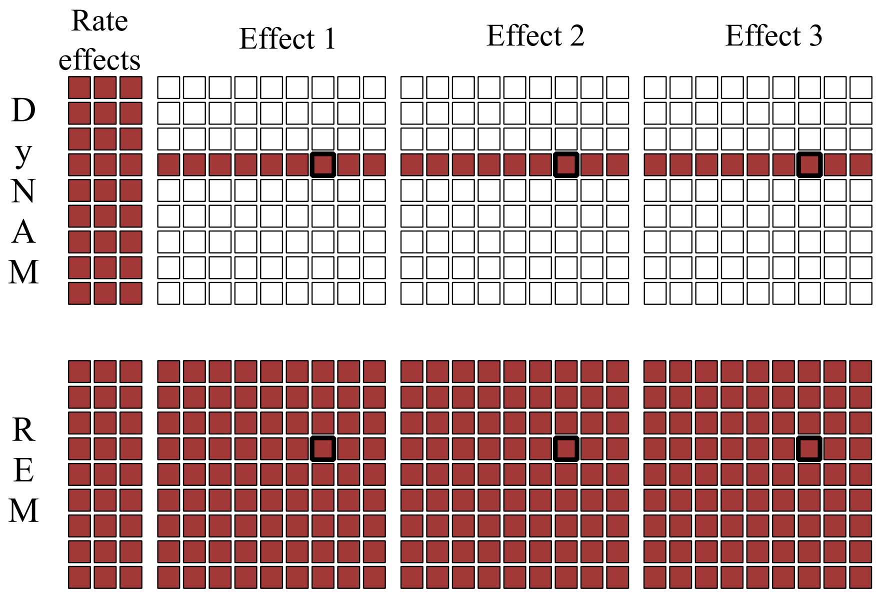

Figure 6: Information used (change of statistics given an event between any pair of actors) for the estimation of DyNAM and REM, respectively. The matrix cells of model statistics of the observed event are marked with a frame.

not explain activity but rather do preexisting network structures. One could now consider extending the model so that individual rates gain a better precision in case of 3:00 a.m. phone calls between old and regularly interacting friends.

Comparison of Computational Complexity

Table 1 shows that the relational event model needs significantly more CPU time (by a factor of about 60). Some differences may stem from a more elaborate implementation of the DyNAM using a Fisher scoring method; however, the major part of the differences is related to the fact that within the REM framework, a higher amount of information is processed. Figure 6 visualizes this difference: in the DyNAM, a number of parameters that do not vary within actors (e.g., outdegree in the friendship network) are used to specify the rates. Of those parameters that vary within actors (e.g., transitive closure of different choice alternatives), only those of the chosen sender are taken into account. In contrast, within the REM rates, all model statistics of all actors are compared for each event. Roughly speaking, this corresponds to a computational complexity that is N times larger, where N is the number of actors (84 in the example). Further complications arise because of higher memory demands. However, those do not affect the reported CPU times in Table 1 because the data set chosen could be processed with the RAM of a single computer. In general, one can conclude that because of a lower computational complexity, DyNAMs are better suited for large data sets with a high number of actors. The estimation routines for both models are parallelized within the `goldfish` package so that this potential drawback of the tie-oriented REM can be overcome and model selection can often be made based on theoretical assumptions.

Biases When Ignoring Right-Censored Intervals

We argue that it is important from a statistical perspective to consider right-censored intervals when modeling relational event sequences and that disregarding them should lead to biased estimates. But what are the biases induced in our example models? Table 2 compares estimates of the full DyNAM rate model (left) and estimates that just take into account the order of events given the process state at actual event times. In both models, the parameters are clearly significant and point into the same direction. However, the absolute values of the two actor effects that relate to the time window effects are larger in the model that ignores right-censoring, whereas the other two effects are very similar. Comparing estimates of exponential family models across models and across data has to be taken with a grain of salt, and we thus refrain from substantively interpreting these differences. It is notable, however, that the absolute change of the parameters is in line with our earlier expectations: when we ignore modeling lagged updates that describe time windows in which actors are potentially more active (one hour in case of our case study), we will overestimate time window–related effects.

The biases are a lot more pronounced when comparing an ordered relational event model that ignores right-censored intervals with a full relational event model as proposed in this article. This comparison is shown in Table 3. Of 15 effects, nine would have been interpreted differently in a model ignoring right-censored intervals, either because the parameter sign or level of significance changed. The interpretation of the differences is not straightforward because the interpretation of each effect in this model is net of all others. The biased effects that relate to structural change are now much closer to the DyNAM receiver choice sub model (see Table 1), which is not concerned with timing but solely models receiver choices once an actor considers making a phone call. The effects that are modeled as DyNAM rates are very different, however, and two of those have a different sign ("outdegreeX"). It is thus unclear how to interpret the estimates of the REM that is ignoring right-censoring, but we can conclude that the induced biases are substantial.

Conclusions

Interpersonal interactions are at the core of major social network theories. However, empirical research to date mostly focuses on the study of stable relations (such as friendship, kinship, and collaboration) rather than on processes of interpersonal interaction (such as contact opportunity, communication, and knowledge exchange). We argued that the availability of new data collection techniques and advances in statistical network models may lead to a paradigm shift in how empirical research addresses questions about the dynamics of social networks by focusing on the dynamic patterns of interpersonal interaction. In an attempt to contribute to this new perspective, this article introduced DyNAMs for the study of directed interaction sequences through time.

DyNAMs are relational event models in the spirit of the seminal work by Butts (2008) but build upon the actor-oriented paradigm of Snijders (1996), a combination of two research lines that was discussed by Stadtfeld and Geyer-Schulz (2011).

Table 2: Comparison of the bias of the DyNAM rate induced by ignoring right-censored intervals.

Effects	(1) $\hat{\theta}$ (s.e.)	(2) $\hat{\theta}'$ (s.e.)	Δ
intercept	$-13.74\dagger$		
	0.05		
egoX recentCallsSent	0.59†	1.14†	0.55
	0.04	0.05	
egoX recentCallsReceived	0.53†	1.29†	0.75
	0.05	0.05	
outdegreeX friendship	0.03†	0.02†	0.01
	0.00	0.00	
outdegreeX callNetwork	0.26†	0.23†	0.03
	0.01	0.01	

Column (1) includes a continuous-time DyNAM rate model, column (2) a sequential model ignoring right-censoring. Standard errors are below estimates.
$* \ p < 0.05$, † $p < 0.01$

DyNAMs for the study of (undirected) coordination network have been introduced in Stadtfeld et al. (2017). We argued that our model fulfills three core criteria that we believe are useful both theoretically and methodologically. First, the actor-oriented nature of the model aligns with a variety of theoretical explanations that consider individual preferences, available interaction opportunities, and, potentially, individuals' perceptions about their position in the social system. Second, the DyNAM is a network model that can take into account how social interactions are embedded in a social system with a complex structure that consists of, for example, various dynamically changing networks and nodal and global covariates. Third, we elaborated on how to study the absolute and relative timing of event sequences and explained under what circumstances models that take into account time windows and exogenous changes return unbiased estimates. This article presented the theoretical foundations of DyNAMs, their mathematical formulation, a numerical estimation method, and a detailed case study in which we compared it to the tie-oriented REM (Butts 2008).

We believe that this is a good time for proposing a new method to enrich the toolbox of social scientists for the analysis of network data, given the constantly growing interest in analyzing relational events over the past decade. Recent examples include the study of communicational dynamics of animals (Tranmer et al. 2015) and emergency responders (Butts 2008), brokerage and receiver choice in communication networks (Quintane and Carnabuci 2016; Stadtfeld, Geyer-Schulz, and Allmendinger 2011), e-mail communication in organizations (Perry and Wolfe 2013), interaction within teams (Leenders, Contractor, and DeChurch 2016), exchange of patients between hospitals (Kitts et al. 2017; Vu et al. 2017), and collaboration on online platforms (Vu et al. 2011). Further, international conflict relations were analyzed in a relational event framework (Lerner et al. 2013).

Table 3: Comparison of the bias of the REM induced by ignoring right-censored intervals.

Effects	(1) $\hat{\gamma}$ (s.e.)	(2) $\hat{\gamma}'$ (s.e.)	Δ
intercept	$-14.78\dagger$		
	0.14		
egoX recentCallsSent	$0.55\dagger$	$0.17\dagger$	0.38
	0.03	0.05	
egoX recentCallsReceived	$-0.27\dagger$	$0.19*$	0.46
	0.05	0.08	
outdegreeX friendship	$-0.04\dagger$	$-0.04\dagger$	0.01
	0.01	0.01	
outdegreeX callNetwork	-0.03	$-0.11\dagger$	0.08
	0.02	0.02	
outdegree callNetwork	$-4.97\dagger$	$-5.44\dagger$	0.47
	0.13	0.14	
outdegree callNetworkPastHour	$0.79\dagger$	$-1.95\dagger$	2.74
	0.11	0.14	
reciprocity callNetwork	$0.38\dagger$	0.09	0.28
	0.11	0.10	
reciprocity callNetworkPastHour	$4.67\dagger$	$2.53\dagger$	2.14
	0.13	0.17	
inPop callNetwork	0.01	$-0.13\dagger$	0.13
	0.03	0.03	
inPop friendship	$-0.06\dagger$	$-0.10\dagger$	0.05
	0.01	0.01	
transitivity callNetwork	$-0.16*$	$0.22\dagger$	0.38
	0.07	0.08	
transitivity friendship	$0.14\dagger$	$0.16\dagger$	0.02
	0.02	0.02	
sameX floor	$-0.98\dagger$	$-0.43\dagger$	0.55
	0.08	0.08	
sameX gradeType	$-0.17*$	$-0.19*$	0.03
	0.08	0.09	
X friendship	$1.52\dagger$	$1.41\dagger$	0.11
	0.12	0.13	

Column (1) includes a continuous-time REM, column (2) a sequential model ignoring right-censoring. Standard errors are below estimates.
$* \ p < 0.05$, $\dagger \ p < 0.01$

A significant part of this article was concerned with contrasting the DyNAM to the tie-oriented REM. We explained the theoretically different assumptions of actor- and tie-oriented network models, building upon the cross-sectional comparison of Block et al. (2016). The DyNAM consists of two sub processes. One models timing of events and actor activity and one the choice of event receivers. In REMs, all parameters in the model contribute to the modeling of timing, sender, and receiver, and thus it has clear advantages over the DyNAM in a likelihood-based model comparison. The fact that REM parameters simultaneously model timing and position of events, however, means that the parameters are typically more difficult to interpret. We showed that ignoring timing of events can potentially lead to significant biases of parameter estimates. In the empirical case study, we could illustrate that interpretation of DyNAMs is rather straightforward, which facilitates the testing of hypotheses that are either related to timing or to the positions of events. A side effect of the DyNAM's separability in two sub processes is a much lower computational complexity in the estimation process and thus a much faster estimation and applicability to larger networks. We conclude that there are no objective criteria that could deem either model superior, but we suggest that researchers should base their choice of the modeling paradigm (actor-oriented DyNAM or tie-oriented REM) on theoretical assumptions in relation to their hypotheses prior to evaluating fit-based criteria.

The development of the DyNAM framework is tightly connected with the development of a software package in R (called `goldfish`), even though in principle researchers may find use in implementing this statistical approach in their own software.

Future extensions of the DyNAM framework and the software could, for example, be concerned with the simultaneous dynamics of one-mode and two-mode interactions, the development of advanced goodness-of-fit criteria, the introduction of advanced techniques for the analysis of timing that go beyond our time window approach, the simultaneous modeling of interpersonal interactions and individual behavioral actions, and the co-evolution of relational events and network ties.

We believe that the study of social interactions is a timely and exciting topic. Interpersonal interactions are at the core of many sociological theories and they are tightly linked to the dynamics of more stable social network relations. We hope that DyNAMs will contribute to a deeper understanding of the role of interactions in the dynamics of complex social networks.

Notes

1 Butts (2008) adds "action type" as a fourth element. This is a natural extension to the most basic case of three elements.

2 Kitts (2014) makes a similar distinction and argues that social network analysis should be concerned with four dimensions of networks: roles, sentiments, interactions, and exchange.

3 Naturally, dependence between events only occurs forward in time.

4 The hazard rate is thus "flat" (Box-Steffensmeier and Jones 2004:22) The process that we describe is Markovian. The flatness of the hazard rate can be linked to the "memorylessness" of the Markov process.

5 For a detailed introduction of exponential event history models we refer to Box-Steffensmeier and Jones (2004:p.22–25).

6 Note that this is contrary to differences between other actor- and tie-based models, in which the model statistics do differ between models. See Block et al. (2016).

7 The DyNAM for undirected events as introduced in Stadtfeld et al. (2017) contains three modeling steps: (i) who will propose a tie, (ii) to whom the tie is proposed, and (iii) whether this proposal is accepted.

8 The software package is available online at www.social-networks.ethz.ch/research/goldfish.html (retrieved 04/2017). The name refers to the (incorrect) assertion that goldfish—like Markov processes—are memoryless. The software has been developed in R (R Core Team 2013). Replication material is available as an online supplement.

9 Stadtfeld and Geyer-Schulz (2011) proposed to combine exponentially weighted effects with a threshold model that preserves the periodwise stability of the process states.

10 We only use a subsample of the available data between the first and the third survey wave; thus, the friendship network is only updated once.

11 Calculated by the inverse of the "empty" rate $\frac{1}{\exp(-13.74)}$. The result is in seconds.

12 Calculated as $\frac{1}{\exp(-13.74+5*0.59+5*0.53+5*0.03+5*0.26)}$

13 $\frac{1}{exp(-13.74+0.59)} / \frac{1}{exp(-13.74)} = \frac{1}{exp(0.59)} = 0.55$

14 $\exp(2.07) = 7.9$.

15 Calculated as $\exp(1.52) = 4.6$.

References

Block, Per. 2015. "Reciprocity, transitivity, and the mysterious three-cycle." *Social Networks* 40:163–173. https://doi.org/10.1016/j.socnet.2014.10.005.

Block, Per, Christoph Stadtfeld, and Tom A. B. Snijders. 2016. "Forms of Dependence: Comparing SAOMs and ERGMs from Basic Principles." *Sociological Methods and Research* E-pub ahead of print. https://doi.org/10.1177/0049124116672680.

Boda, Zsófia and Balinth Néray. 2015. "Inter-ethnic friendship and negative ties in secondary school." *Social Networks* 43:57–72. https://doi.org/10.1016/j.socnet.2015.03.004.

Box-Steffensmeier, Janet M. and Bradford S. Jones. 2004. *Event History Modeling: A Guide for Social Scientists*. Cambridge: Cambridge University Press.

Burt, Ronald S. 2004. "Structural Holes and Good Ideas." *American Journal of Sociology* 110:349–399. https://doi.org/10.1086/421787.

Butts, Carter T. 2008. "A relational event framework for social action." *Sociological Methodology* 38:155–200.

Cox, David R. and David Oakes. 1984. *Analysis of survival data*, volume 21. CRC Press.

Cramer, J.S. 2003. *Logit Models from economics and other fields*. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511615412.

Deuflhard, Peter. 2004. *Newton Methods for Nonlinear Problems*. Number 35 in Springer Series in Computational Mathematics. Springer-Verlag Berlin Heidelberg.

Emerson, Richard M. 1976. "Social Exchange Theory." *Annual Review of Sociology* 2:335–362. https://doi.org/10.1146/annurev.so.02.080176.002003.

Feld, Scott L. 1997. "Structural embeddedness and stability of interpersonal relations." *Social Networks* 19:91–95. https://doi.org/10.1016/S0378-8733(96)00293-6.

Festinger, Leon, Stanley Schachter, and Kurt Back. 1950. *Social Pressures in Informal Groups*. Stanford University Press. 1974 preprint.

Friedkin, Noah E. 1998. *A Structural Theory of Social Influence*. Cambridge University Press. https://doi.org/10.1017/CBO9780511527524.

Hedström, Peter and Peter Bearman. 2009. "What is analytical sociology all about? An introductory essay." *The Oxford Handbook of Analytical Sociology* pp. 3–24.

Heider, Fritz. 1958. *The Psychology of Interpersonal Relation*. John Wiley & Sons. https://doi.org/10.1037/10628-000.

Hollway, James and Johan Koskinen. 2016. "Multilevel embeddedness: The case of the global fisheries governance complex." *Social Networks* 44:281–294. https://doi.org/10.1016/j.socnet.2015.03.001.

Huston, Ted L. and George Levinger. 1978. "Interpersonal attraction and relationships." *Annual Review of Psychology* 29:115–156. https://doi.org/10.1146/annurev.ps.29.020178.000555.

Kadushin, Charles. 2012. *Understanding Social Networks. Theories, Concepts and Findings*. New York: Oxford University Press.

Kitts, James A. 2014. "Beyond Networks in Structural Theories of Exchange: Promises from Computational Social Science." *Advances in Group Processes* 31:263–298. https://doi.org/10.1108/S0882-614520140000031007.

Kitts, James A., Alessandro Lomi, Daniele Mascia, Francesca Pallotti, and Eric Quintane. 2017. "Investigating the Temporal Dynamics of Interorganizational Reciprocity: Patient Exchange Among Italian Hospitals." *American Journal of Sociology* Forthcoming.

Lazer, David, Alex (Sandy) Pentland, Lada Adamic, Sinan Aral, Albert Laszlo Barabasi, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, Tony Jebara, Gary King, Michael Macy, Deb Roy, and Marshall Van Alstyne. 2009. "Computational Social Science." *Science* 323:721–723. https://doi.org/10.1126/science.1167742.

Leenders, Roger, Noshir S. Contractor, and Leslie A. DeChurch. 2016. "Once upon a time: Understanding team processes as relational event networks." *Organizational Psychology Review* 6:92–115. https://doi.org/10.1177/2041386615578312.

Lerner, Juergen, Margit Bussmann, Tom A. B. Snijders, and Ulrik Brandes. 2013. "Modeling frequency and type of interaction in event networks." *Corvinus Journal of Sociology and Social Policy* 4:3–32.

Lomi, Alessandro and Christoph Stadtfeld. 2014. "Social Networks and Social Settings: Developing a Coevolutionary View." *Kölner Zeitschrift für Soziologie und Sozialpsychologie* 66:395–415. https://doi.org/10.1007/s11577-014-0271-8.

Lospinoso, Josh. 2012. *Statistical Models for Social Network Dynamics*. Ph.D. thesis, Magdalen College, University of Oxford.

Lusher, Dean, Johan Koskinen, and Garry Robins (eds.). 2013. *Exponential Random Graph Models for Social Networks*. Cambridge University Press.

Madan, Anmol, Manuel Cebrian, Sai Moturu, Katayoun Farrahi, and Alex Sandy Pentland. 2012. "Sensing the 'Health State' of a Community." *IEEE Pervasive Computing* 11:36–45.

Marcum, Christopher S. and Carter T. Butts. 2015. "Creating Sequence Statistics for Egocentric Relational Events Models using informR."." *Journal of Statistical Software* 64:1–34.

McFadden, Daniel. 1974. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers in Econometrics*, edited by Paul Zarembka, chapter 4, pp. 105–142. New York: Academic Press Inc.

Pál, Judit, Christoph Stadtfeld, André Grow, and K. Takács. 2016. "Status Perceptions Matter: Understanding Disliking Among Adolescents." *Journal of Research on Adolescence* 26:805–818. https://doi.org/10.1111/jora.12231.

Perry, Patrick O. and Patrick J. Wolfe. 2013. "Point process modelling for directed interaction networks." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 75:821–849. https://doi.org/10.1111/rssb.12013.

Quintane, Eric and Gianluca Carnabuci. 2016. "How Do Brokers Broker? Tertius Gaudens, Tertius Iungens, and the Temporality of Structural Holes." *Organization Science* pp. 1343–1360. https://doi.org/10.1287/orsc.2016.1091.

R Core Team. 2013. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.

Ripley, Ruth M., Tom A. B. Snijders, Zsófia Boda, András Vörös, and Paulina Preciado Lopez. 2016. *Manual for SIENA 4.0.* Nuffield College and Department of Statistics, University of Oxford. Version: December 6, 2013.

Rivera, Mark T., Sara B. Soderstrom, and Brian Uzzi. 2010. "Dynamics of Dyads in Social Networks: Assortative, Relational, and Proximity Mechanisms." *Annual Review of Sociology* 36:91–115. https://doi.org/10.1146/annurev.soc.34.040507.134743.

Schaefer, David R., Jimi Adams, and Steven A. Haas. 2013. "Social Networks and Smoking: Exploring the Effects of Peer Influence and Smoker Popularity Through Simulations." *Health Education and Behavior* 40:24–32. https://doi.org/10.1177/1090198113493091.

Snijders, Tom A. B. 1996. "Stochastic actor-oriented models for network change." *The Journal of Mathematical Sociology* 21:149–172. https://doi.org/10.1080/0022250X.1996.9990178.

Snijders, Tom A. B. 2001. "The Statistical Evaluation of Social Network Dynamics." In *Sociological Methodology*, edited by M.E. Sobel and Becker M.P., volume 31, pp. 361–395. Boston and London: Basil Blackwell.

Snijders, Tom A. B. 2005. "Models for Longitudinal Network Data." In *Models and methods in social network analysis*, edited by Peter J. Carrington, John Scott, and Stanley Wasserman, volume 27 of *Structural Analysis in the Social Sciences*, chapter 11, pp. 215–247. New York: Cambridge University Press.

Snijders, Tom A. B., Alessandro Lomi, and Vanina Jasmine Torlo. 2013. "A model for the multiplex dynamics of two-mode and one-mode networks, with an application to employment preference, friendship, and advice." *Social Networks* 35:265–276. https://doi.org/10.1016/j.socnet.2012.05.005.

Snijders, Tom A. B., Philippa E. Pattison, Garry L. Robins, and Mark S. Handcock. 2006. "New specifications for exponential random graph models." *Sociological Methodology* 36:99–153. https://doi.org/10.1111/j.1467-9531.2006.00176.x.

Snijders, T. A. B. and Christian. Steglich. 2015. "Representing Micro–Macro Linkages by Actor-based Dynamic Network Models." *Sociological Methods & Research* 44:222–271. https://doi.org/10.1177/0049124113494573.

Snijders, Tom A. B., Gerhard G. van de Bunt, and Christian E. G. Steglich. 2010. "Introduction to Stochastic Actor-based Models for Network Dynamics." *Social Networks* 32:44– 60. https://doi.org/10.1016/j.socnet.2009.02.004.

Stadtfeld, Christoph. 2012. *Events in Social Networks. A Stochastic Actor-oriented Framework for Dynamic Event Processes in Social Networks*. KIT Scientific Publishing. http://digbib.ubka.uni-karlsruhe.de/volltexte/1000025407.

Stadtfeld, Christoph and Andreas Geyer-Schulz. 2011. "Analyzing Event Stream Dynamics in Two Mode Networks: An Exploratory Analysis of Private Communication in a Question and Answer Community." *Social Networks* 33:258–272. https://doi.org/10.1016/j.socnet.2011.07.004.

Stadtfeld, Christoph, Andreas Geyer-Schulz, and Otto Allmendinger. 2011. "The Influence of Distance, Time, and Communication Network Structures on the Choice of Communication Partners." In *Proceedings of the 2011 IEEE International Conference on Social Computing (SocialCom)*.

Stadtfeld, Christoph, James Hollway, and Per Block. 2017. "Dynamic Network Actor Models: Investigating Coordination Ties through Time." *Sociological Methodology* 47.

Tranmer, Mark, Christopher S. Marcum, F. Blake Morton, Darren P. Croft, and Selvino R. de Kort. 2015. "Using the relational event model (REM) to investigate the temporal dynamics of animal social networks." *Animal Behaviour* 101:99–105. https://doi.org/10.1016/j.anbehav.2014.12.005.

Vörös, András and Tom A. B. Snijders. 2017. "Cluster analysis of multiplex networks: Defining composite network measures." *Social Networks* 49:93–112. https://doi.org/10.1016/j.socnet.2017.01.002.

Vu, Duy, Alessandro Lomi, Daniele Mascia, and Francesca Pallotti. 2017. "Relational event models for longitudinal network data with an application to interhospital patient transfers." *Statistics in Medicine* https://doi.org/10.1002/sim.7247.

Vu, Duy Q., Arthur U. Asuncion, David R. Hunter, and Padhraic Smyth. 2011. "Continuous-Time Regression Models for Longitudinal Networks." *Advances in Neural Information Processing Systems* 2492–2500.

Wang, Peng, Garry Robins, Pip Pattison, and Emmanuel Lazega. 2013. "Exponential random graph models for multilevel networks." *Social Networks* 35:96–115. https://doi.org/10.1016/j.socnet.2013.01.004.

Acknowledgements: Useful feedback and comments that considerably improved this article were graciously provided by James Hollway, by members of the social network research group in Zürich, by participants of the Swiss Networks Workshop in Zürich, and by the 9th International Conference on Social Science Methodology (RC33) in Leicester. Discussions with Alessandro Lomi and Viviana Amati contributed to the formulation of the idea of network mechanisms that operate on different time scales—they refer to this idea as "process time."

Christoph Stadtfeld: Department of Humanities, Social and Political Sciences, ETH Zürich. E-mail: c.stadtfeld@ethz.ch.

Per Block: Department of Humanities, Social and Political Sciences, ETH Zürich. E-mail: per.block@gess.ethz.ch.