UvA-DARE (Digital Academic Repository)

Networks of action and events over time: a multilevel discrete-time event history model for longitudinal network data

de Nooy, W.

DOI

[10.1016/j.socnet.2010.09.003](https://doi.org/10.1016/j.socnet.2010.09.003)

Publication date

2011

Published in

Social Networks

[Link to publication](#)

Citation for published version (APA):

de Nooy, W. (2011). Networks of action and events over time: a multilevel discrete-time event history model for longitudinal network data. *Social Networks*, 33(1), 31-40.
<https://doi.org/10.1016/j.socnet.2010.09.003>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Networks of action and events over time

A multilevel discrete-time event history model for longitudinal network data

DRAFT – PLEASE CITE THE PUBLISHED ARTICLE (*SOCIAL NETWORKS*, 33 (2011) 1, P. 31-40).

Wouter de Nooy*

The Amsterdam School of Communication Research ASCoR, Department of Communication, University of Amsterdam, Kloveniersburgwal 48, 1012 CX Amsterdam, The Netherlands.

E-mail: w.denooy@uva.nl, tel. +31(0)20.525.2123, fax +31(0)20.525.3681.

Abstract

Longitudinal network data recording the moment at which ties appear, change, or disappear are increasingly available. Event history models can be used to analyze the dynamics of time-stamped network data. This paper adapts the discrete-time event history model to social network data. A discrete-time event history model can easily incorporate a multilevel design and time-varying covariates. A multilevel design is needed to account for dependencies among ties and vertices, which should not be ignored in a small longitudinal network. Time-varying covariates are required to analyze network effects, that is, the impact of previous ties. In addition, a discrete-time event history model handles constraints on who can act or who can be acted upon in a straightforward way. The model can be estimated with multilevel logistic regression analysis, which is illustrated by an application to book reviews, so network evolution can be analyzed with a fairly standard statistical tool.

Keywords: network dynamics; discrete-time event history models; multilevel logistic regression analysis; longitudinal social networks; book reviewing.

In the analysis of longitudinal social networks, relatively little attention has been paid to the timing of relational events, namely the moments at which ties appear, change, or disappear. Initially, this kind of data was not available because longitudinal data were collected by means of repeatedly administered surveys (Katz and Proctor 1959: 318; Runger and Wasserman 1979: 144). As a consequence, models for network dynamics

* Acknowledgements: Thanks to Mark Tranmer (University of Manchester) for his advice and suggestions on previous versions of this paper, which were presented at the Sunbelt XXVIII International Social Network Conference 2008 (St. Petersburg, FLA) and the 7th International Amsterdam Multilevel Conference 2009 (Amsterdam). Comments from anonymous reviewers and the co-editor were very useful to improve readability, fix flaws in the reasoning, and extend the literature covered.

have been developed for panel data, including recent advances such as the actor-based statistical models for social dynamics implemented in SIENA (Snijders 2005). Time-stamped relational event data, however, are increasingly available from digitalized archives, electronic devices recording the exact time of online and offline communications (discussion lists, tracking devices), experimental designs (Callander and Plott 2005; Corten and Buskens 2010), and so forth. They offer the opportunity to test hypotheses on the timing of network change: Who acts upon whom at what time?

In social, medical, and engineering research, the peculiarities of event data gave rise to a family of statistical models – event history models, also called survival models, duration models, or failure-time models – for testing statements on whether and when events occur and which covariates affect the hazard of event occurrence. Section 1 discusses the ways in which event history models have been applied to social network data, noting that a discrete-time or multilevel event history model has not yet been used for predicting tie formation even though it handles complications of network data in an efficient way. The discrete-time event history model is specified in Section 2. Section 3 shows that empirical constraints on who can act or who can be acted upon are easily included in this model. Section 4 presents an application of the model to empirical data. The concluding section compares the discrete-time event history model to other models for time-stamped network data; it discusses extensions to the discrete-time event history model and implications for the collection of network data.

1 Event history models applied to social networks

Event history models aim to explain whether and when events occur. These models typically assume a process with a known starting time in which actors are at risk of experiencing an event until the event happens or observation stops. In applications of event history analysis to social network data, the event to be explained is either an abrupt change of the state an actor is in or the occurrence of a new interaction involving a pair of actors. The first type of application tests a social influence model (Friedkin and Johnsen 1990; Robins et al. 2001b) because the social network explains why an actor changes his mind, behaviour, or situation. Peers are assumed to influence a social actor through their network ties. The second type of application exemplifies the social selection model, which aims to predict the networking behaviour of social actors: Why do they engage in one tie and not in another (Robins et al. 2001a)?

Event history models have mainly been applied to social network data from the perspective of the social influence model. Marsden and Podolny (1990) introduced event history designs for social influence models in the social network analysis community. Event history models predicting the timing of an action or event from characteristics of the person's or organization's network position have frequently been used although it should be noted that the network position sometimes merely represents whether one is married or the number of children one has. The impact of the size and composition of a person's social network on health or mortality is a popular topic (Adams et al. 2002; Bygren et al. 1996; Gustafsson et al. 1998; Kang and Bloom 1993; Litwin and Shiovitz-Ezra 2006; Patterson et al. 1996; Payette et al. 2000; Samuelsson and Dehlin 1993; Trovato 1998; Villingshoj et al. 2006) and so is the impact of exposure to adopters of an innovation through social or geographical network ties (Bogart 2007; Bohman 2006;

Chaves 1996; Davis and Stout 1992; Edling and Sandell 2001; Lipp and Krempel 2001; McKeown 1994; Mintrom and Vergari 1998; Soule 1997; Strang 1991; Strang and Tuma 1993; Van den Bulte and Lillien 2001). The relevance of network ties for the hazard of finding a job has also been studied (Bernasco et al. 1998; Brandt 2006; Yakubovich and Kozina 2000) as well as the hazard of leaving a field or organization (Clarysse et al. 1996; McPherson et al. 1992; Mossholder et al. 2005; Sutton and Chaves 2004).

In contrast, this paper focuses on the use of event history models in explaining social selection, that is, the genesis of new network links. Applications are rare here. Some work addresses the propensity to be involved in ties rather than tie formation itself. Kim and Higgins (2007), for example, analyzed the hazard of biotechnical firms forming an alliance, but they did not predict which firms had become allies. Similarly, Tsai predicted the hazard of intra-organizational linkages (Tsai 2000). Robinson and Smith-Lovin investigated speaker turns assuming that the rest of the group is always the addressee, so they did not need to look at specific links in the network (Robinson and Smith-Lovin 2001).

Applications that explain link formation have to deal with two technical complications. First, there are dependencies between lines (pairs) and the two vertices or social actors that constitute the pair, namely the line's tail and head. Vertices appear repeatedly as tail or head in pairs and each pair joins a tail and a head. A multilevel model may adequately account for these dependencies. Second, network context must be included as a covariate if typical network effects such as reciprocation of a previous line are tested but network context varies with time. Each new line changes network context at the next time. If network context is limited to the lines that appeared in a limited preceding time period—a retrospective sliding window approach (Moody et al. 2005: 1212)—each disappearing line also changes the network context. Similarly, if a decay function is used to weigh network context by the length of time passed, every tick of the clock changes the value of the network context covariates.

All applications of event history models analyzing link formation that I have found conceptualize time as a continuous phenomenon and none uses time-varying covariates or a multilevel model. The earliest example of a continuous-time event history model explaining tie formation is Krempel's study of a longitudinal network among university freshmen (Krempel 1990). Krempel used a parametric continuous-time event history model, but he neither corrected for dependencies between pairs (lines) and vertices (actors), nor did he include time-varying covariates in the analysis.

Kossinets and Watts (2006) used the continuous-time semi-parametric Cox regression model to analyze determinants of transitive closure within pairs that were linked by a path of length two that was not yet closed by a direct arc. The appearance of such a path marked the start of an episode, the occurrence of an arc closing the path ended it. They avoided the dependency problem by sampling pairs from a very large recorded network, so they had independent observations. Unfortunately, this approach cannot be applied to longitudinal data for a small network. Kossinets and Watts did not include time-varying covariates although local network context could have changed during the episode and other properties of network context such as reciprocity may have offered alternative explanations for the formation of the arc. Kossinets and Watts's approach was also used for analyzing networks among criminals (Hu et al. 2009).

The final example is the continuous-time piece-wise constant rate model for a small network followed over time developed by Butts (2008) to explain the length of waiting times between successive actions, namely responder radio communications after the World Trade Center disaster. Here, dependencies caused by repeated observations of the same persons across different pairs are solved by adding dummies for all persons minus one. This is not a parsimonious solution, so it is difficult to detect interesting results, which was acknowledged in the paper (Butts 2008: 185-186). Butts' model predicted the duration between successive actions without limit to the retrospective window and without weighting, so the network context was constant during each interval; no time-varying covariates were needed here. In terms of event history models, each interval was treated as a separate episode because the process time clock was reset at each event; after each event all people were at risk again. As a consequence, the estimated overall piece-wise hazard function, which concatenates all piece-wise constant hazards, is different from hazard functions in event history models, which do not reset the clock during the process. Butts' method was extended by Brandes et al. (2009) to account for competing risks.

2 *A multilevel discrete-time event history model for network data*

This section presents a multilevel discrete-time event history model for investigating link formation in a longitudinal network. The model requires time-stamped network data: the time at which a network tie starts, changes, or ends must be recorded. I propose a discrete-time event history model rather than a continuous-time model because the book reviews analyzed here (Section 4) represent a process in discrete time; critics publish reviews typically at a fixed weekday. At the other days of the week, the critics are not at risk of publishing a review. Tied time observations occur a lot, that is, reviews appearing on the same day or week, so a discrete-time model is appropriate.

Discrete-time processes are found in several social networks due to external constraints, e.g., institutionally arranged moments at which representatives can vote. In addition, continuous-time processes must sometimes be analyzed with discrete-time models if the processes are measured with coarse time indicators, which occurs often in archival data as well as data from memory elicited with an Event Oriented Observation Plan (Blossfeld et al. 1989: 24) or the Life History Calendar method (Freedman et al. 1988). Note that event history models do not necessarily require very fine-grained time measurement. The timescale used in data measurement is the timescale at which event occurrence is analyzed, so it should first and foremost match the timescale of the process under investigation.

With continuous-time processes, one may either group the observations into subperiods and apply a discrete-time event history model or use a continuous-time event history model. The piecewise-constant exponential model is a continuous-time alternative to the model proposed here, which is also capable of handling multiple levels and time-varying covariates.[†]

The hazard function is the quantity of major interest in event history models. Discrete-time hazard is defined as the conditional probability of an event happening to individual i

[†] This was pointed out to me by an anonymous reviewer, which I gratefully acknowledge.

in time period t given that no event has occurred in an earlier time period (Singer and Willett 2003: 330). Following the notation of Steele (2008), discrete-time hazard h_{ti} for individual i at time t is:

$$h_{ti} = \Pr[y_{ti} \neq 0 | y_{si} = 0, s < t] \quad (1)$$

with $y = 0$ indicating the absence of an event. In the data analyzed here, the appearance of a book review in a newspaper or magazine is the event to be explained.

In discrete-time event history analysis, data are stored in a person-period format. A person-period data set contains one case for each individual i for each time period t at which the individual is at risk of experiencing the event. If the model is restricted to explaining the occurrence ($y = 1$) versus absence ($y = 0$) of an event, the event history model for a single episode is:

$$\text{logit}(h_{ti}) = \log\left(\frac{h_{ti}}{1 - h_{ti}}\right) = \alpha(t) + \beta^T \mathbf{x}_{ti} \quad (2)$$

with y as the dependent variable, $\alpha(t)$ a function of duration time, e.g., a set of time dummies representing the baseline hazard function, β a vector of effect parameters, and $\mathbf{x}_{ti}$ a vector of covariates. This model can be estimated with logistic regression analysis of the person-period data set.

Note that time-varying covariates can be included just as easily as time-constant covariates because the data set contains a case for each time period in which a person is at risk. A case may contain the current value of a time-varying covariate or the fixed value of a time-constant covariate. The logit link function specifies a proportional odds model, but other link functions can also be used, e.g., a complementary log-log link function leads to a discrete-time approximation to a proportional hazards model.

If individuals may experience events more than once, there are multiple episodes in which individuals are at risk. In the case of book reviews, the publication of each new book starts and defines an episode. Each critic is at risk to review the new book in the time periods within the episode. The person-period data set is expanded to contain a case for each combination of person (i , e.g., a critic), period (t , e.g., a week), and episode (j , e.g., a period started by the publication of a new book). Then, hazards within the same individual are likely to be related among episodes, e.g., some critics review relatively quickly. Steele (2008: 14-15) describes a multilevel model accounting for dependencies at the individual level:

$$\text{logit}(h_{tji}) = \log\left(\frac{h_{tji}}{1 - h_{tji}}\right) = \alpha(t) + \beta^T \mathbf{x}_{tji} + u_i. \quad (3)$$

This model contains a random effect $u_i \sim \mathbf{N}(0, \sigma^2)$, which captures the effect of individual-specific unobserved variables.

Note that episodes may overlap in time, e.g., a critic is at risk of reviewing several recently published books in the same week. The model does not directly take into account the choice among available books but it evaluates each book separately. It merely assumes that books with more favourable properties will be picked sooner. The risks are not competing in the sense that reviewing one book excludes or affects the possibility that the critic reviews another book; the other book can be reviewed at the same or a later time. In the example of literary criticism, think of a review evaluating two or more books. In general, the model assumes that an event may happen at each of its opportunities and that

the hazard of one event to happen is independent of the hazard of another event that did happen or could have happened at the same time.

An important advantage of the multilevel model is that the random factor absorbs unmeasured heterogeneity among individuals with respect to the hazard of experiencing an event; it is a shared-frailty model. Unmeasured heterogeneity is a problem because the more resistant observations, e.g., slow reviewers, remain in the risk set longer (Box-Steffensmeier and Jones 2004: 142-148; Steele 2003).

If one wants to apply this model to network data, one must note that the (ordered) pair is the basic unit of analysis, not the individual person. In the example analyzed here, a critic reviews a book by an author. Instead of a person-period data setup, a pair-period data file is needed with one case for each pair at risk in each time period and episode. With the pair as the basic unit of analysis, more dependencies must be taken into account because each pair involves two individuals, the person who is acting (the tail f , e.g., the critic) and the person who is acted upon (the head g , e.g., the reviewed author) in a directed network. The random effect of the individual i must be replaced by two cross-classified random effects, one for tails (v_f) and one for heads (w_g), which represent propensities to act or be acted upon. In the example, the diligence of a critic is captured by the random effect of the tail and the random effect of the head shows the popularity of an author as a target for reviews. Cross-classified random effects are described in Goldstein (1995: Ch. 8) and this design was introduced in the analysis of network data with the p2 model (Van Duijn et al. 2004).

If the same person can appear as tail and head, e.g., some authors also publish reviews, one may consider additional constraints, viz., that the random effects are correlated for tails and heads that refer to the same person. This is called a multiple roles model (Snijders and Bosker 1999: 161-162). In a directed network, acting is usually very different from being acted upon, so random effects of tails and heads do not have to be correlated. Distinguishing between the actor (e.g., a critic) and the person acted upon (e.g., an author), I have hitherto assumed that the network is directed. However, the model also applies to undirected networks; just substitute the ordered pair by an unordered pair.

Finally, I propose to add a random effect for episodes (u_j). If an event outside the network defines the episode, that is, starts the process clock, this event may have characteristics that affect the hazard function. In the example, the new book that starts the episode may have characteristics that favour a quick review, e.g., because it was published by an important publishing house. Effects of episodes' characteristics are best modelled as a random factor. Equation 4 presents the final model, with each random effect u_j (for episodes), v_f (for tails), and w_g (for heads) assumed to be normally distributed with mean zero. Note that pairs are cross-classified by tail, head, and episode because they belong to the tail, head, and episode simultaneously.

$$\mathbf{logit}(h_{tfg}) = \log\left(\frac{h_{tfg}}{1 - h_{tfg}}\right) = \alpha(t) + \beta^T \mathbf{x}_{tfg} + u_j + v_f + w_g \quad (4)$$

This model may include all types of predictors that are commonly used in social network analysis: attributes of the tail (critic) and head (author), separately or in combination (similarity, ranking), a tie on another network relation, and characteristics of local network structure (e.g., network of previously published reviews) such as reciprocity, transitivity, and balance. Local network covariates are calculated in the manner proposed

by Wasserman and Pattison (1996): A network context covariate counts theoretically interesting subnetworks that would be created by a new arc from the selected tail to the selected head. For example, each path of length two turns into a transitive triad if a new arc is added from the starting point to the endpoint of the path. More paths of length two from a selected tail to a selected head yield a higher score for the transitivity covariate. They also raise the odds of creating a new arc from the tail to the head over not creating it, if transitivity guides link formation. A tendency toward transitivity, then, shows up as a positive parameter estimate for the transitivity covariate in the logistic regression analysis. Likewise, a negative parameter estimate shows a tendency to avoid transitivity.

Table 1 - Example of a pair-period data matrix.

Occasion_ID (<i>i</i>)	Episode_ID (<i>j</i>)	tail_ID (<i>f</i>)	head_ID (<i>g</i>)	Period (<i>t</i>)	Event (<i>y</i>)	TCC (x_{ifg})	TVC (x_{ijfg})
1	1	6	3	1	0	0	8
2	1	6	3	2	0	0	8
3	1	6	3	3	1	0	9
4	1	9	3	1	0	1	8
5	1	9	3	2	1	1	8
6	1	13	3	1	0	1	8
7	1	13	3	2	0	1	8
8	1	13	3	3	0	1	9
9	1	13	3	4	0	1	10
10	1	13	3	5	0	1	10
11	2	6	3	1	1	0	10
12	2	15	3	1	0	1	10
13	2	15	3	2	0	1	11
14	2	15	3	3	1	1	11
15	3	3	5	1	1	0	0
16	3	6	5	1	0	1	0
17	3	6	5	2	1	1	1

Table 1 shows an example of a small pair-period data matrix. It features three books (1, 2, and 3) identified in the *Episode_ID* (*j*) column. Books 1 and 2 are written by Author 3 (column *head_ID*, *g*) while Book 3 is written by Author 5. Critic 6 (column *tail_ID*, *f*) reviews (a ‘1’ in column *Event*, *y*) Book 1 in the third week after its publication (column *Period*, *t*), Book 2 in the first week after its publication, and Book 3 in the second week. Author 3 also reviews books, e.g., Book 3 by Author 5 in the first week after its publication. Critic 13 is right-censored in the episode of Book 1: after five weeks, she has still not published a review on it.

The column headed *TCC* (x_{ifg}) shows a time-constant covariate, viz., whether the critic and author have the same sex (‘1’) or different sex (‘2’). This dyadic covariate tests the effects of homophily with respect to sex in the event history model. In the example, Author 5 and Critic 6 are male and the other critics and authors are female. The final column, headed *TVC* (x_{ijfg}) contains a time-varying covariate, namely the popularity of the author measured as the number of previous reviews of his current and previous books. This is a simple characteristic of local network context. Note that Author 3 had received 8 reviews of previous books when Book 1 appeared, whereas Book 3 is Author 5’s first book so he had no previous reviews. Popularity counts increment in the week after the publication of a review.

Note that local network structure covariates are calculated from links that precede the link to be predicted. The circular dependencies that prohibit the application of regression models to cross-sectional network data (Wasserman and Robins 2005) do not apply here

because a link may depend on a preceding link but not the other way around. This is, at least, the assumption of the model. As a consequence, modelling reciprocity requires neither a multilevel structure as in the Social Relations Model for cross-sectional network data (Snijders and Kenny 1999) nor a bivariate (van Duijn et al. 2004) or multinomial regression model (Zijlstra et al. 2009) as in the p2 model. In this regard, the time ordering of network links simplifies the p2 model. In addition, it allows the inclusion of network effects beyond the pair, such as transitivity.

3 *Constraints on eligible pairs*

Event history models typically presuppose a meaningful starting time after which people are at risk of experiencing an event, for example, acting toward someone else. The occasion starting the process clock may have characteristics that affect whether and when people act, but it may also restrict the set of people that are in the position to act. In the example of book reviewing, the reviewed author is fixed by the book that starts the episode as it can only be this book's author. Note that literary books, which are analyzed here, usually have just one author.

Peculiarities of the empirical situation may lead to complicated institutional constraints on the eligible pairs, that is, who can act and who can be acted upon at particular times. Critics may, for example, be active only during part of the period studied because they enter the field late or retire from it during the period. Naturally, they are only at risk of publishing reviews during the period in which they are active. Table 1, for example, does not contain entries for Critic 15 in the episode of Book 1, which signifies that this critic was not active at the time the book appeared. Note another constraint: authors are not allowed to review their own books, so there are no entries with Critic 3 for Book 1 and Book 2. These are examples of exogenous circumstances limiting the set of actors that are at risk (comp. Butts 2008). Constraints on the units that are at risk within a time interval are called discontinuous risks in event history models and the situation in which some units are never at risk is known as a split population (Box-Steffensmeier and Jones 2004: 148-154). Constraints on eligible pairs are implemented in the discrete-time event history model simply by including only pairs that can interact in the pair-period data set.

A literary book is reviewed in a newspaper or magazine at most once. As soon as a periodical publishes a review of a book, none of the periodical's critics is at risk of publishing a review of this book in this periodical any longer, so the pair-period data set must not contain cases for them. This is an example of an endogenous constraint, caused by an event included in the data set. Again, the constraint is implemented by including the right cases in the pair-period data set.

Social relations are often constrained in terms of the temporal availability of actors and restrictions on who can act or who can be acted upon. Command and organizational structures restrict who can order whom, heterosexual and homosexual partner choices pose obvious restrictions on eligible partners, each job opening determines the starting time of the application process, but it also fixes the organization receiving the applications in vacancy chains (White 1970), and so forth. The pair-period data format can handle these constraints.

4 *Application*

To demonstrate the model, I apply it to the publication of literary book reviews in newspapers and magazines. My research question concerns the time at which reviews appear: Can the model predict whether a critic reviews a new book and, if so, how soon after its publication? This question addresses the news value of a literary book and its author as well as the preferences and diligence of critics. I test four substantive hypotheses:

1. The number of evaluations of an author's work published in the previous 2 years increases the hazard of the author's new book to be reviewed.

An author who has recently received a lot of attention is "hot" and a new book by this author has higher news value. This is the popularity effect in social network models.

2. An author's commercial success, measured as having reprints of books, affects the hazard that a new book by this author is reviewed.

There are competing views on the direction of this effect. In general, commercial success increases the news value of a product or person because a larger part of the media's audience is expected to be interested. In contrast, Bourdieu's field theory predicts that commercial success is depreciated in a cultural field such as the literary field (Bourdieu 1983). In his view, the odds of a review are expected to be lower for commercially successful literary authors.

3. The first book by an author—the debut—has a higher hazard of being reviewed than subsequent books.

According to a previous study (Van Rees and Vermunt 1996), literary critics are more likely to review an author's first book than his or her later books.

4. If a critic evaluated a previous book by an author in the preceding two years, the hazard increases that this critic reviews a new book by the same author.

Critics tend to follow authors whom they review because they generally like the work they review and tend to patronize their favourite authors. This is the conformity or inertia effect in social network modelling.

4.1 *Data*

The data were collected and analyzed by De Nooy, see his publications for details about the data collection and previous results (De Nooy 1991; 1999; 2008). The original data set contains all reviews and interviews among 40 literary authors and critics who appear most frequently in a debate on contemporary literature in the Netherlands, 1970-1980. In addition, it contains all other reviews of the authors' works in this period, so the data set can be said to represent the public debate among the selected authors and critics within the wider context of book reviewing in that period.

The current analysis investigates only the reviews of new books of literary prose published by the 28 authors in this data set. The aim is to predict the occurrence and timing of reviews. Each review yields an arc from a critic (tail) to a reviewed author (head). The publication of a new book defines an episode and starts the process clock. A multiple-episodes event history model is needed because there are 95 books in the data set, ranging from one to ten books per author.

Each review is dated by the day it was published, which I aggregated to weeks because newspapers usually have weekly book sections and magazines appear only once a week. Due to the publication rhythm, it makes more sense to predict the week in which a review appears than the precise day. The date of a new book's availability to critics is set to the week before the appearance of the first review. The pair-period data set contains cases (reviewing opportunities) for time periods up to and including 20 weeks. Hardly any reviews were published after week 20 and if they were, other reasons for reviewing the book than its appearance had to be reckoned with.

The local network context of a review is constructed from reviews and interviews published in the preceding 104 weeks (two years). A decay function—weighting according to the length of time passed—is not used. Note that local network context takes into account all evaluations in the original network and not just the reviews that are predicted in the event history analysis. The distinction between network context and lines to be explained has implications, which are discussed in the final section.

There are several constraints on the set of eligible critic–author pairs. First, each episode (new book) has just one author who can be reviewed. The one book with multiple authors in the data set is treated as a different book for each author. A second constraint applies to the critics. The analysis is restricted to the main critics of the periodicals, that is, the critics who reviewed new books of literary prose in a periodical regularly, because it is difficult to predict reviews by incidental reviewers. A total of 31 critics remain in the analysis, who are at risk of publishing a review in a periodical only during the period in which they regularly published reviews in this periodical. A final institutional constraint is posed by the practice that a book is reviewed only once in a periodical. If a book has been reviewed in a periodical, other critics working for the periodical cannot publish another review.

The final pair-period data set contains 26,185 cases or reviewing opportunities. This number is considerably less than the product of the number of episodes (95 books), the number of critics (31), and the number of weeks within each episode (20). First, not all critics are active during the entire period. Second, critics who publish a review on a book are no longer eligible because a critic reviews a book only once, so no cases are created for the remaining weeks for them. In addition, their peers working for the same periodical are right-censored at the time of the review because a periodical reviews a book only once.

Event occurrence, that is, the appearance of a review, is rare, occurring in no more than 1.5 percent of all pair-period cases. It should be noted that critics often do not review books; 78 percent of all critic-book combinations never result in a review. As a consequence, the event history results reveal the hazard of getting a book review to a larger extent than the speed at which a book is reviewed in this application. Therefore, the interpretation focuses on *whether* a book is reviewed rather than *when* it is reviewed.

Table 2 - Description of the variables.

	Minimum	Maximum	Average	SD	TVC avg. range	TVC max. range
Event occurrence (y)	0	1	0.015	-	-	-
Activity (sqrt)	0.00	5.00	2.10	1.05	0.24	2.14
Popularity (ln)	0.00	4.28	2.36	0.73	0.81	2.77
Commercial success: reprints	0	1	.44	-	.06	1
Debut: an author's first (literary) book	0	1	.15	-	-	-
Conformity: previous review	0	1	.16	-	.03	1
Transitivity: 2-path from tail to head	0	1	.26	-	.11	1

Table 2 presents a description of the variables used in the analysis. Variables measuring activity and transitivity were added to the variables mentioned in the hypothesis, because they represent common network covariates that should be controlled for. The activity variable captures the number of reviews published by the critic in the preceding 24 months. This count is skewed, so the square root is taken. This is a time-varying covariate because a critic may publish new reviews on other authors during the time he is at risk of publishing a review on the book defining the episode. In addition, old reviews may drop out of the retrospective sliding window as time passes, also changing the critic's activity score. The last two columns describe the variability of the time-varying covariates. For each critic within each episode the range of the activity covariate (among others) is calculated, which represents the maximum change in the covariate for the critic. The average range over all critic-episode combinations is 0.24 and the maximum range occurring is 2.14. These results show that there is quite some variation in a critic's activity covariate scores within an episode.

The natural logarithm is taken of the number of reviews received by the author in the preceding 24 months to obtain a measure of popularity that is not too skewed. The variation of popularity scores within each episode-author combination is even larger than the variation of the activity covariate over time. Commercial success, which is another characteristic of the author, is a dummy variable indicating whether an author's book had reprints. Summarizing over all cases in the data set, 44 percent of all reviewing occasions concern authors with at least one reprint at the time. In about 6 percent of all episode-author combinations, a first reprint occurred during the episode, changing the author's commercial success score from zero to one.

Debut is a time-constant variable within an episode indicating whether the reviewed book is the author's first literary book. It applies to 15 percent of the episodes. The last two rows of Table 2 present properties of network context as dyadic covariates, that is, characteristics of the pair. Conformity—the occurrence of a review by the critic on the author in the preceding 24 months—occurs in 16 percent of the cases in the data set, while transitivity—the presence of a two-step path from the critic to the author in the network of evaluations in the preceding 24 months—occurs slightly more often. The variation over time within an episode is largest for transitivity: about 11 percent of all episode-pair combinations experienced a change from no previous path to a path or vice versa during an episode.

Table 3 – Pearson correlations among covariates.

	Activity	Popularity	Commercial success	Debut	Conformity
Popularity	-.03				
Commercial success	.01	.43			
Debut	.00	-.28	-.38		
Conformity	.32	.34	.14	-.12	
Transitivity	.35	.25	.13	-.09	.19

The correlations among the covariates is low to moderate (Table 3; note that four out of six covariates are dichotomies). Activity and popularity correlate with transitivity, indicating that it is sensible to control for transitivity while determining the effects of activity and popularity. This also holds for conformity. In addition, popularity, commercial success, and debut are correlated at moderate strength, so we may expect that the addition or removal of one of these effects changes the effects of the other two covariates. Debut is negatively associated with popularity and commercial success because it signals the first book by an author, which precludes reviews (popularity) and reprints (commercial success) of previous books.

4.2 Results

The model was estimated with multilevel logistic regression analysis using MCMC estimation as implemented in the software package MLwiN (Browne 2004) with a burn-in length of 500 iterations[‡] and 50,000 runs for the estimation. As a preliminary step (not shown in Table 4), variance components models were analyzed. Significant variances were found at the critic (tail) and book (episode) level with a substantial reduction of the Deviance Information Criterion (DIC), which is a goodness-of-fit index generalized from Akaike's Information Criterion (Spiegelhalter et al. 2002). Some critics and books had overall higher odds of publishing reviews or of being reviewed. The addition of the reviewed book's author as a random factor deteriorated the DIC value. There may be no substantial variation among authors that is not covered by the variation among their books, which is not surprising because a majority of authors have just one or few books in the data set. In the analyses, none of the estimated models turned out to have a significant random effect for authors on top of the random effect for books, so the effect is not reported in Table 4.

Model 1 (Table 4) presents estimates of the baseline hazard function. In comparison to linear, quadratic, and cubic functions of the time passed since the start of the episode, dummy variables perform better in terms of DIC. The categorization presented in Table 4 is the most parsimonious with the best fit. After the first week of the episode, the log-odds of publishing a review drop in each subsequent period. The probability that a book is reviewed is estimated to be low in the first week; about one out of twenty reviewing opportunities results in a review ($\text{antilogit}(-3.10) = 0.043$). The probability drops to about 1 out of 60 in week 2 to 4, and less than one out of 300 after week 11.

[‡] This number of burn-in runs seems to suffice. Model 2 estimated with 5000 burn-in runs produced parameter estimates that differed less than 0.01 from the ones presented here.

Table 4 – Posterior estimates of fixed and random effects on the log-odds of being reviewed in a given week.

	Model 1		Model 2	
	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>
<i>Fixed Part</i>				
Constant	-3.10	0.20***	-2.90	0.22***
Time (reference category: week 1)				
week 2 – 4	-0.99	0.14***	-1.09	0.15***
week 5 – 8	-1.36	0.15***	-1.48	0.16***
week 9 – 11	-1.88	0.20***	-2.03	0.22***
week 12 – 20	-2.73	0.19***	-2.90	0.20***
Activity (sqrt, grand mean centred)			0.30	0.08***
Popularity (ln, grand mean centred)			0.32	0.11**
Commercial success: reprints			-0.37	0.19 ^o
Debut: an author's first (literary) book			0.47	0.26 ^o
Conformity: previous review			0.27	0.16 ^o
Transitivity			-0.22	0.15
<i>Random Part</i>				
$\sigma^2_{critics}$	0.63	0.22*	0.36	0.15*
σ^2_{books}	0.41	0.11***	0.43	0.12***
# critics (tails)	31		31	
# books (episodes)	95		95	
# cases (evaluation opportunities)	26,185		26,185	
-2*loglikelihood				
DIC	3448.54		3430.16	

Dependent variable: yes/no review. MCMC estimation, 50,000 runs.

*** Wald (joint) chi square test significant at 0.1% (two-tailed test).

** Wald chi square test significant at 1% (two-tailed test).

* Wald chi square test significant at 5% (two-tailed test).

^o Wald chi square test significant at 5% (one-tailed test).

In Model 2, all covariates are included. The activity of the critic, which controls for differences among frequent versus infrequent reviewers, is statistically significant and positive, showing that critics who paid a lot of attention to the authors in the previous 2 years have a higher risk of reviewing any new book. This covariate reduces the random-effect variance among the critics. The transitivity effect was also included for reasons of control but this effect is not statistically significant. The network tends toward a 2-mode network with one set of actors specializing in reviewing and the other set predominantly being reviewed; transitivity effects are less likely here.

Commercial success and debut are nearly statistically significant at the 5 percent level (two-sided test) provided that the model controls for popularity. Without popularity, they have insignificant and negligible effects. If commercial success or debut is removed, the effect of the other covariate becomes statistically significant at the 5 percent level. In Section 4.1, it was noted that the three covariates are linked: popularity and commercial success go together and they are less likely to accompany an author's first book. Thus it makes sense that we need to control for one covariate to get significant results for the other and that commercial success and debut are partly confounded. In comparison to other books by authors with few previous reviews, debuts have a higher hazard of being reviewed and commercially successful books have a lower hazard compared to books by authors that have just as many previous reviews. The model's DIC value improves substantially with any of these covariates included, so all three covariates remain in the

model. Finally, the effect of conformity is significant at the 5 percent level only in a one-sided test.

The results for the substantive hypotheses are as follows. The first hypothesis is confirmed, there is a positive effect of an author's popularity among literary critics in the preceding two years on the odds of a review for a new book by the author. The result for the hypothesis about commercial success corroborates Bourdieu's notion of an economic world reversed rather than the news value of commercial success; an author's previous commercial success reduces the odds of getting a review (more quickly), but the effect appears only if the author's popularity is controlled for. In other words, commercial success decreases the odds of a (quick) review compared to authors who are equally popular among critics, but less commercially successful.

The first book by an author (debut) has a higher hazard of being reviewed than subsequent books (predicted by hypothesis 3). This result was also found in a previous study with a different design (Van Rees and Vermunt 1996). Finally, the hazard that a critic reviews a book increases if this critic also reviewed a previous book by the same author. Critics have a tendency to follow literary authors.

Although the results are substantively interesting, the main goal of the application is to illustrate how the multilevel discrete-time event history model can be applied to network data. Estimation is straightforward and the analysis yields results that are in line with the expectations.

5 *Conclusion and discussion*

A multilevel event history model is a flexible instrument for analyzing the occurrence and timing of relational events, such as actions and interactions within a relatively small social group. The aim is to explain whether and when actors act toward one another. How soon and under which circumstances do network ties appear, change, or disappear? A range of questions regarding network evolution, notably the speed at which networks develop, can be answered with a multilevel event history model. The model proposed here applies to data measured in discrete time, but it can also be used for data measured in continuous time if time is recoded into discrete intervals. In addition, a multilevel regression model is also available for continuous-time data, viz., a Poisson regression model.

The event-history model assumes that actors are at risk to act after the onset of an initial event, which starts an episode. This implies that they may react to the initial event. In the example analyzed here, the initial event is the publication of a new book and the risk is that a critic publishes a review of it. The time between the initial event and action is modelled. The model can handle multiple episodes, both concurrent episodes, e.g., several new books appearing in the same period, and non-overlapping episodes. Previous actions that appear within or before the episode constitute the network context for new action. In a multi-episode design, then, the model may take into account two types of predictors: characteristics of the episode, e.g., some types of books may be reviewed more quickly, and characteristics of the actors and the network of their interactions as it is at the onset of the initial event and as it develops subsequently. This is different from the model proposed by Butts (2008), which analyzes the waiting time between successive actions. Butts' model assumes that new action is primarily triggered by the last preceding action. This is similar to an event-history model in which each new action closes the preceding

episode and starts a new one in which typically all actors are at risk again. The choice between the two models should be motivated by the relevance of the two assumptions to the data at hand: are there initial events that actors respond to or is the process driven mainly by each last action?

Another difference between Butts' approach and the model proposed here is the demarcation of the retrospective window capturing the network context in which action is embedded. Butts uses all previous action within the group since the onset of the WTC disaster. This approach assumes that actions before the disaster are not relevant, which is a plausible assumption. In contrast, my application uses a sliding window, taking into account the reviews published in the preceding 24 months. The substantive motivation here is that a new book is not released in a social vacuum and critics are aware of reviews of previous books. As a consequence, actions before the initial event are relevant.

The cumulative approach of Butts has practical advantages because one can create the relevant network context covariates even for the very first observation (though their values are zero) and the researcher needs not decide on the length of the retrospective window. In contrast, a sliding window requires a decision on the appropriate length and network covariates can be calculated only for observations in the data set after a full time span of the sliding window, so the first observations cannot be predicted. Both approaches may use a decay function attributing more weight to recent actions. Again, a choice by the researcher is needed but little is known about the salience of previous ties and the decay rate at which previous ties tend to be forgotten or ignored. Empirical work is called for on this topic. However, results tend to be robust with sliding windows of different length (Kossinets and Watts 2009, p. 416) as long as they include sufficient observations. This may also be true for specifications of the decay function. After all, the general trend will be that recent events are memorized better.

A general problem with statistically modelling of social networks is the sparseness of data: most pairs are unconnected rather than connected. In a discrete-time event history model, time aggravates this problem because pairs are unconnected for every single moment until the event takes place or they are right-censored, so the data are even more sparse than in a cross-sectional network. As a consequence, the model works better, e.g., it is easier to estimate, if there are constraints on who can act toward whom at what time. Constraints eliminate instances without action, reducing the proportion of zeros in the dependent variable. The model is particularly suited for experimental designs that can be regulated by the researchers, e.g., game-theoretic experiments, and naturally occurring data in settings with institutional constraints. Book review data are an example of the latter because reviews by selected critics appear in designated periodicals at scheduled moments. These types of constraints are usually found in processes occurring in the mass media but many other social processes are regulated as well, for example, decision-making within organizations or politics, recruitment of board members, and so on. Note that institutional constraints are also responsible for the co-occurrence of events at set times, which favours a discrete-time approach over a continuous-time model.

The model proposed in this paper works better if it is applied to a rather small set of actors that interact frequently because there is a better balance between present and absent lines in such a network. As a consequence, however, actors appear repeatedly as senders (tails) and/or receivers (heads) of action, so we must correct for the dependencies among observations within actors. This can be done effectively with random effects in a

multilevel model. If the ties, however, are drawn at random from a very large network, which is the case in Kossinets & Watts' application (Kossinets and Watts 2006), repeated observations on actors are very unlikely to occur, so a multilevel model is not needed. A random factor in a multilevel model presupposes that the entities may be regarded as a random sample, so it is perfectly OK to sample from a small network and use a multilevel model to account for dependencies among observations. Especially in small but dense networks for which the measurement of ties is time-consuming, sampling may be a good strategy.

There is a similarity between Kossinets & Watts' model and the one proposed in this paper: the analysis is restricted to a subset of the vertices and lines in the network. In the Kossinets & Watts design, tie occurrence is estimated for pairs linked by a two-step path. The intermediate vertex on the path is just part of the network context in their analysis. This also applies to the research reported here: the book reviews by the selected critics on the selected authors (the focal actors) are analyzed within a wider context of book reviews by other critics. This has an important implication for data collection and the network boundary problem. In an actor-based approach using local network context as a predictor, one can collect all ties relevant to the network context in a limited number of snowballing steps. Recall that network context covariates are counts of subnetworks including the selected tail and head. The transitivity covariate, for example, counts all two-step paths from the selected tail to the selected head in the network of previous links. Assuming that the subnetworks counted by the covariate are connected, the maximum distance between one of the focal actors and any other vertex in the subnetworks is equal to the number of snowballing rounds required to collect data. To test for transitivity, for example, all neighbours (step 1 in snowballing) of the focal actors must be included and all lines linking them to the focal actors must be added to the data set. If network context includes only vertices at short distance from the focal vertices, data collection requires few snowballing steps. Full coverage of the network is not necessary hence the vexing problem of network boundaries is avoided.

Estimating the discrete-time event history model with a multilevel logistic regression analysis of a pair-period data set, the major hurdle that needs to be taken is the construction of the data set. As argued above, the model is expected to work especially well on data with constraints on who can act toward whom at what time. As a consequence, the creation of cases for pairs eligible for tie formation may not be trivial. In addition, time-varying covariates must be calculated from network context. These processes, however, can be computerized; a Windows application for the creation of a pair-period data set from a time-stamped network data file is available at the author's web site (<http://home.medewerker.uva.nl/w.denooy/>).

The multilevel discrete-time event history model is a General Linear (Mixed) Model, so it can easily be extended. Instead of a dichotomous dependent variable, which distinguishes between the presence and absence of a link, a variable with three or more categories can be used contrasting the absence of a link to the presence of two or more types of links. This is a competing risks event history model, predicting both the occurrence and the quality of the link (cf. Brandes et al. 2009). It can be estimated with a multilevel multinomial regression model. This model can be used, for example, to predict whether a book is reviewed and, if so, whether it is reviewed positively or negatively. Further extensions to multivariate event history models have been developed, for

example, the competing risks with multiple states model used by Steele et al. (2004). Perhaps it is even possible to estimate multivariate models that combine the social selection and social influence models, explaining both link formation and changing attributes of the actors. An example here could be the co-evolution of the network of book reviews and commercial success of authors among readers.

Simplifications of the model are also possible. If research interest focuses on explaining whether a tie is formed during a fixed period but timing is not important, e.g., whether two people at a university start to have e-mail contact during a period of over 200 days (Kossinets and Watts 2009) or whether two scientists start to co-author papers in a particular year (O'Madadhain et al. 2005), one may aggregate the pair-period data set keeping just one case per pair over the selected period with a variable indicating the presence or absence of a tie. Another simplified model takes for granted the historical fact that a tie occurred, e.g., a review was published, and focuses on explaining the nature or contents of the action, e.g., whether the review was positive or negative (De Nooy 2008). In this model, the pair-period data set needs to contain a case only for each published review and a variable telling whether it was positive or negative. The model promises to be a versatile instrument for answering a wide range of research questions on time-stamped longitudinal network data.

References

- Adams, A. M., Madhavan, S., Simon, D., 2002. Women's social networks and child survival in Mali. *Social Science Medicine* 54, 165-178.
- Bernasco, W., De Graaf, P. M., Ultee, W. C., 1998. Coupled careers - Effects of spouse's resources on occupational attainment in the Netherlands. *European Sociological Review* 14, 15-31.
- Blossfeld, H.-P., Hamerle, A., Mayer, K. U., 1989. *Event History Analysis. Statistical Theory and Application in the Social Sciences*. Lawrence Erlbaum, Hillsdale, NJ.
- Bogart, D., 2007. Neighbors, networks, and the development of transport systems: Explaining the diffusion of turnpike trusts in eighteenth-century England. *Journal of Urban Economics* 61, 238-262.
- Bohman, L., 2006. Economic action and interfirm relations: Diffusion of stock repurchases on the Stockholm Stock Exchange 2000-2003. *European Sociological Review* 22, 383-396.
- Bourdieu, P., 1983. The field of cultural production, or: the economic world reversed. *Poetics* 12, 311-356.
- Box-Steffensmeier, J. M., Jones, B. S., 2004. *Event History Modeling. A Guide for Social Scientists*. Cambridge University Press, Cambridge.
- Brandes, U., Lerner, J., Snijders, T. A. B., 2009. Networks Evolving Step by Step: Statistical Analysis of Dyadic Event Data, 2009 International Conference Advances in Social Network Analysis and Mining (ASONAM 2009). IEEE Computer Society, pp. 200-205.

- Brandt, M., 2006. Finding a job via social contacts. *Kölner Zeitschrift für Soziologie und Sozialpsychologie* 58, 468-488.
- Browne, W. J., 2004. MCMC Estimation in MLwiN. Centre for Multilevel Modelling, London.
- Butts, C. T., 2008. A relational event framework for social action. *Sociological Methodology* 38, 155-200.
- Bygren, L. O., Konlaan, B. B., Johansson, S. E., 1996. Attendance at cultural events, reading books or periodicals, and making music or singing in a choir as determinants for survival: Swedish interview survey of living conditions. *British Medical Journal* 313, 1577-1580.
- Callander, S., Plott, C. R., 2005. Principles of network development and evolution: an experimental study. *Journal of Public Economics* 89, 1469-1495.
- Chaves, M., 1996. Ordaining women: The diffusion of an organizational innovation. *American Journal of Sociology* 101, 840-873.
- Clarysse, B., Debackere, K., Rappa, M. A., 1996. Modeling the persistence of organizations in an emerging field: The case of hepatitis C. *Research Policy* 25, 671-687.
- Corten, R., Buskens, V., 2010. Co-evolution of conventions and networks: An experimental study. *Social Networks*, 32, 4-15.
- Davis, G. F., Stout, S. K., 1992. Organization theory and the market for corporate-control: A dynamic analysis of the characteristics of large takeover targets, 1980-1990. *Administrative Science Quarterly* 37, 605-633.
- De Nooy, W., 1991. Social networks and classification in literature. *Poetics* 20, 507-537.
- , 1999. A literary playground. *Literary criticism and balance theory. Poetics* 26, 385-404.
- , 2008. Signs over time: Statistical and visual analysis of a longitudinal signed network. *Journal of Social Structure* 9.
- Edling, C., Sandell, R., 2001. Social influence and corporate behaviour: A case study of interdependent decision-making in Sweden's publicly traded firms. *European Sociological Review* 17, 389-399.
- Freedman, D., Thomson, A., Cambum, D., Alwin, D., Young-DeMarco, L., 1988. The Life History Calendar: A technique for collecting retrospective data. *Sociological Methodology* 38, 37-68.
- Friedkin, N. E., Johnsen, E. C., 1990. Social Influence and Opinions. *Journal of Mathematical Sociology* 15, 193-205.
- Goldstein, H., 1995. *Multilevel Statistical Models*. London: Arnold.
- Gustafsson, T. M., Isacson, D. G. L., Thorslund, M., 1998. Mortality in elderly men and women in a Swedish municipality. *Age and Ageing* 27, 585-593.
- Hu, D. N., Kaza, S., Chen, H. C., 2009. Identifying significant facilitators of dark network evolution. *Journal of the American Society for Information Science and Technology* 60, 655-665.

- Kang, S. H., Bloom, J. R., 1993. Social support and cancer screening among older Black-Americans. *Journal of the National Cancer Institute* 85, 737-742.
- Katz, L., Proctor, C. H., 1959. The configuration of interpersonal relations in a group as a time-dependent stochastic process. *Psychometrika* 24, 317-327.
- Kim, J. W., Higgins, M. C., 2007. Where do alliances come from? The effects of upper echelons on alliance formation. *Research Policy* 36, 499-514.
- Kossinets, G., Watts, D. J., 2006. Empirical analysis of an evolving social network. *Science* 311, 88-90.
- 2009. Origins of Homophily in an Evolving Social Network. *American Journal of Sociology* 115, 405-450.
- Krempel, L., 1990. Interpersonal structure and contact. Empirical evidence from the analysis of a series of social networks in time and contact as a process. An analysis with tools of 'Event History' Analysis. In: Weesie, J., Flap, H. (Eds.), *Social Networks Through Time*. ISOR, Utrecht, pp. 65-90.
- Lipp, C., Krempel, L., 2001. Petitions and the social context of political mobilization in the revolution of 1848/49: A microhistorical actor-centred network analysis. *International Review of Social History* 46, 151-169.
- Litwin, H., Shiovitz-Ezra, S., 2006. Network type and mortality risk in later life. *Gerontologist* 46, 735-743.
- Marsden, P. V., Podolny, J. M., 1990. Dynamic analysis of network diffusion processes. In: Weesie, J., Flap, H. (Eds.), *Social Networks Through Time* ISOR, Utrecht, pp. 197-214.
- McKeown, T. J., 1994. The epidemiology of corporate PAC formation, 1975-84. *Journal of Economic Behavior Organization* 24, 153-168.
- McPherson, J. M., Popielarz, P. A., Drobnic, S., 1992. Social networks and organizational dynamics. *American Sociological Review* 57, 153-170.
- Mintrom, M., Vergari, S., 1998. Policy networks and innovation diffusion: The case of state education reforms. *Journal of Politics* 60, 126-148.
- Moody, J., McFarland, D., Bender-deMoll, S., 2005. Dynamic network visualization. *American Journal of Sociology* 110, 1206-1241.
- Mossholder, K. W., Settoon, R. P., Henagan, S. C., 2005. A relational perspective on turnover: Examining structural, attitudinal, and behavioral predictors. *Academy of Management Journal* 48, 607-618.
- O'Madadhain, J., Hutchins, J., Smyth, P., 2005. Prediction and Ranking Algorithms for Event-Based Network Data. *SIGKDD Explorations Newsletter* 7, 23-30.
- Patterson, T. L., Shaw, W. S., Semple, S. J., Cherner, M., McCutchan, J. A., Atkinson, J. H., Grant, I., Nannis, E., 1996. Relationship of psychosocial factors to HIV disease progression. *Annals of Behavioral Medicine* 18, 30-39.

- Payette, H., Coulombe, C., Boutier, V., Gray-Donald, K., 2000. Nutrition risk factors for institutionalization in a free-living functionally dependent elderly population. *Journal of Clinical Epidemiology* 53, 579-587.
- Robins, G., Elliott, P., Pattison, P., 2001a. Network models for social selection processes. *Social Networks* 23, 1-30.
- Robins, G., Pattison, P., Elliott, P., 2001b. Network models for social influence processes. *Psychometrika* 66, 161-189.
- Robinson, D. T., Smith-Lovin, L., 2001. Getting a laugh: Gender, status, and humor in task discussions. *Social Forces* 80, 123-158.
- Runger, G., Wasserman, S., 1979. Longitudinal analysis of friendship networks. *Social Networks* 2, 143-154.
- Samuelsson, G., Dehlin, O., 1993. Family network and mortality: Survival chances through the life-span of an entire age cohort. *International Journal of Aging Human Development* 37, 277-295.
- Singer, J. D., Willett, J. B., 2003. *Applied Longitudinal Data Analysis: Modeling Change and Event Occurrence*. Oxford University Press, New York.
- Snijders, T. A. B., 2005. Models for longitudinal network data. In: Carrington, P.J., Scott J., Wasserman, S. (Eds.), *Models and Methods in Social Network Analysis*. Cambridge University Press, Cambridge, pp. 215-247.
- Snijders, T. A. B., Bosker, R. J., 1999. *Multilevel Analysis: An Introduction to Basic and Advanced Multilevel Modeling*. SAGE , London.
- Snijders, T. A. B., Kenny, D. A., 1999. The social relations model for family data: A multilevel approach. *Personal Relationships* 6, 471-486.
- Soule, S. A., 1997. The student divestment movement in the United States and tactical diffusion: The shantytown protest. *Social Forces* 75, 855-882.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., Linde, A. v. d., 2002. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society, Series B (Statistical Methodology)* 64, 583-640.
- Steele, F., 2003. A discrete-time multilevel mixture model for event history data with long-term survivors, with an application to an analysis of contraceptive sterilization in Bangladesh. *Lifetime Data Analysis* 9, 155-174.
- 2008. Multilevel models for longitudinal data. *Journal of the Royal Statistical Society Series a-Statistics in Society* 171, 5-19.
- Steele, F., Goldstein, H., Browne, W., 2004. A general multilevel multistate competing risks model for event history data, with an application to a study of contraceptive use dynamics. *Statistical Modelling: An International Journal* 4, 145-159.
- Strang, D., 1991. Adding Social Structure to Diffusion Models: An Event History Framework. *Sociological Methods Research* 19, 324-353.
- Strang, D., Tuma, N. B., 1993. Spatial and temporal heterogeneity in diffusion. *American Journal of Sociology* 99, 614-639.

- Sutton, J. R., Chaves, M., 2004. Explaining schism in American protestant denominations, 1890-1990. *Journal for the Scientific Study of Religion* 43, 171-190.
- Trovato, F., 1998. Nativity, marital status and mortality in Canada. *Canadian Review of Sociology and Anthropology-Revue Canadienne de Sociologie et d'Anthropologie* 35, 65-91.
- Tsai, W. P., 2000. Social capital, strategic relatedness and the formation of intraorganizational linkages. *Strategic Management Journal* 21, 925-939.
- Van den Bulte, C., Lillien, G. L., 2001. Medical innovation revisited: Social contagion versus marketing effort. *American Journal of Sociology* 106, 1409-1435.
- Van Duijn, M. A. J., Snijders, T. A. B., Zijlstra, B. J. H., 2004. p(2): A random effects model with covariates for directed graphs. *Statistica Neerlandica* 58, 234-254.
- Van Rees, C. J., Vermunt, J., 1996. Event history analysis of authors' reputation: Effects of critics' attention on debutants' careers. *Poetics* 23, 317-333.
- Villingshoj, M., Ross, L., Thomsen, B. L., Johansen, C., 2006. Does marital status and altered contact with the social network predict colorectal cancer survival? *European Journal of Cancer* 42, 3022-3027.
- Wasserman, S., Pattison, P., 1996. Logit models and logistic regressions for social networks: 1. An introduction to Markov graphs and p*. *Psychometrika* 61, 401-425.
- Wasserman, S., Robins, G., 2005. An introduction to random graphs, dependence graphs, and p*. In: Carrington, P.J., Scott J., Wasserman, S. (Eds.), *Models and Methods in Social Network Analysis*. Cambridge University Press, Cambridge, pp. 148-161.
- White, H. C., 1970. *Chains of Opportunity: System Models of Mobility in Organizations*. Harvard University Press, Cambridge, Massachusetts.
- Yakubovich, V., Kozina, I., 2000. The changing significance of ties: An exploration of the hiring channels in the Russian transitional labor market. *International Sociology* 15, 479-500.
- Zijlstra, B. J. H., van Duijn, M. A. J., Snijders, T. A. B., 2009. MCMC estimation for the p(2) network regression model with crossed random effects. *British Journal of Mathematical Statistical Psychology* 62, 143-166.