Tilburg University

Modeling the evolution of interaction behavior in social networks

Mulder, Joris; Leenders, Roger

Published in:
Chaos, Solitons & Fractals

DOI:
[10.1016/j.chaos.2018.11.027](https://doi.org/10.1016/j.chaos.2018.11.027)

Publication date:
2019

Document Version
Peer reviewed version

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
Mulder, J., & Leenders, R. (2019). Modeling the evolution of interaction behavior in social networks: A dynamic relational event approach for real-time analysis. *Chaos, Solitons & Fractals*, 119, 73-85.
<https://doi.org/10.1016/j.chaos.2018.11.027>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Modeling the evolution of interaction behavior in social networks: a dynamic relational event approach for real-time analysis

Joris Mulder

Department of Methodology and Statistics & Jheronimus Academy of Data Science, Tilburg University, Warandelaan 2, 5037AB Tilburg, The Netherlands

Roger Th.A.J. Leenders (corresponding author)

Department of Organization Studies & Jheronimus Academy of Data Science, Tilburg University, Warandelaan 2, 5037AB Tilburg, The Netherlands

r.t.a.j.leenders@tilburguniversity.edu

POSTPRINT, ACCEPTED FOR PUBLICATION IN

Chaos, Solitons, & Fractals

<https://www.journals.elsevier.com/chaos-solitons-and-fractals>

Modeling the evolution of interaction behavior in social networks: a dynamic relational event approach for real-time analysis

1. Introduction

There has been an increasing interest in understanding how social, biological, or information networks evolve [1-15]. An observation of the network at a specific time t represents a state of the network from which a researcher can meaningfully calculate network density, centralization, clustering, et cetera. The study of the dynamics of such networks is typically based on a (small) number of snapshots of the network. These snapshots are then used to create models to estimate transition probabilities between states and to derive statistics that follow from these transition rates or that can, themselves, drive these transitions. Statistical models for analyzing such networks have become mathematically and computationally advanced and flexible to understand complex dynamic relationships between actors [16-20].

In the present study, we will focus on a specific type of networks: networks of event streams. Network streams are also known as time-ordered networks [21]. These networks are not (necessarily) characterized by stable edges, but consist of edges that are constantly activated and terminated in real time. Examples of such networks include networks of email messaging, networks of ball passing during a football match, networks of information sharing among police agencies, networks of violent interaction among gangs, networks of ants interacting with each other through their antennae, or networks of chimpanzees grooming each other. Such networks are driven by a constant flow of events where none of these events by itself characterizes the network state at any point in time. This is why the statistical models mentioned above are not suitable to understand how these dynamic networks evolve.

The goal of this paper is to present a methodology for analyzing relational event streams to improve understanding of when, how, and why social interaction behavior changes over time. The methodology builds on the relational event modeling framework of Butts [22], which we extend in several ways. As a first extension, we propose a moving window technique to investigate how drivers of relational events (e.g., reciprocity or nodal characteristics) change over time. Because interaction behavior often changes over time, the best fitting model to explain interaction behavior will often also vary over time. To get a better understanding of which model captures the data best at different points in time and to see how statistical evidence between competing statistical models changes, we propose to use an approximate Bayes factor. Bayes factors translate to posterior probabilities that quantify how plausible each model is given a set of competing models at a certain point in time. This approach results in a simple summary statistic of which model best captures interaction behavior at different time points, and how much better it does at fitting the data than competing models.

The network in the current study consists of employees of a consultancy firm who send emails about innovation projects and ideas about innovation activities to each other over the span of an entire year. At the global level, these messages create an evolving global network structure of email exchange that reflects how employees exchange information by email within this organization. This global structure is driven by the activity of the various nodes (i.e., employees) within the network, as each employee has individual control over when and to whom to send their messages. As a result, the global dynamics of the email network result from the local temporal dynamics among the dyads in the network. Vice versa, the global structure can affect the local dynamics in turn. Hence, these networks form

self-organizing systems of interacting local and global dynamics [23-30]. This topic fits especially with the ideas of “complexity matching” and “management of small teams,” which are two of the core topics in this special issue [31]. In line with the theme of the special issue, the objective of the current paper is to propose a statistical approach to analyze how interactions at a local (i.e., dyadic) level shape interaction dynamics at the global level (i.e., at the level of the firm), and vice versa. A full-fledged study of the dynamics of the email network in this firm is beyond the scope of the paper; rather we provide an insight into some of the drivers of the interaction and illustrate how our approach can be applied to study temporal dynamics in networks of relational event streams.

We organize the paper as follows. In Section 2 the relational event model is described together with a moving window to capture network dynamics. Next, the usage of the model is illustrated using empirical data of email messages in Section 3. In Section 4 it is explained how to compute statistical evidence using the Bayes factor and posterior model probabilities. Section 5 then shows statistical evidence between competing models changes over time in the email network. Finally, the paper ends with some concluding remarks in Section 6.

2. Relational event modeling

The type of network we model is built up by sequences of interactions, called *relational events*, between a sender and a receiver at observed points in time. For each email sent, the event includes who is the sender, who is (are) the receiver(s), and at exactly what time the message was sent. Such events constitute a network of directional ties.

This tuple [sender, receiver, time] can easily be extended by including additional characteristics of the event—the content type of the message, its length, its sentiment, et cetera—but here we will limit ourselves to the tuple [sender, receiver, time]. With current technological developments, it is quite easy to collect sequences of these relational events. Much interaction in modern society occurs through communication technology (e.g., email) leaving easily harvestable digital traces about senders, receivers, and timing. Because these data contain information about relational events in continuous time, such data can potentially tell us how fast/slow teams operate, why and when it speeds up or slows down, how the past affects the future, and how (quickly) social order evolves.

Our statistical approach is an extension of the Relational Event Model (REM) [4, 22, 32], which in itself is an extension of a survival model with time-varying covariates [33-36]. In the REM framework, the time until the next relational event is modeled using an exponential distribution where the rate parameter, denoted by λ , is the sum of the rate parameters defined over all possible directed pairs of possible senders and receivers, i.e., $\sum_{s', r'} \lambda(s', r', t)$. The probability that the next event, after the one that occurred at time t , occurs between sender s and receiver r is then given by $\frac{\lambda(s, r, t)}{\sum_{s', r'} \lambda(s', r', t)}$, which

follows a multinomial distribution. The rate parameter of every directed dyad (sender s , receiver r) at time t depends on endogenous variables and exogenous variables using a log linear function and the current time. The endogenous variables summarize the information of the past event stream up to t . For an email network, endogenous variables could for instance include the number of messages a particular person has received until t (reflecting whether the person is a popular receiver), the number of messages the person sent until t (reflecting whether a person is an active sender), the proportion of messages from s to r that are subsequently forwarded to person v (which might make it tempting for s to skip sending to r in the future and send to v directly). Exogenous variables summarize the external characteristics that affect network dynamics. Examples include the sender's hierarchical level (e.g., low, medium, high), a person's location (e.g., which building), a person's job

type (e.g., consultant, auditor, support staff), gender (male, female). Exogenous variables can be vertex-specific (e.g., a person's hierarchical level) or dyad-specific (e.g., the difference in hierarchical level between a sender and a receiver). An overview of variables that can drive the rates at which events occur in event networks is given in [4].

For example, if the sending rate of emails within a given directed dyad is assumed to depend on 1) the hierarchical difference between sender and receiver (in certain working cultures communication mainly goes either bottom-up or top-down), 2) whether sender and receiver work in the same field (two employees who are both consultants may have a stronger tendency to share information than a consultant and a person working in tax), and 3) whether the sender received information from the receiver in the past (reciprocity), the rate function would be modeled as:

$$\log \lambda(s, r, t) = x_{\text{hierarchical.diff}}(s, r) \times \beta_{\text{hierarchical.diff}} + x_{\text{same.field}}(s, r) \times \beta_{\text{same.field}} + x_{\text{messages.send}}(r, s, t) \times \beta_{\text{reciprocity}}$$

where $x_{\text{hierarchical.diff}}(s, r)$ denotes the hierarchical difference between sender s and receiver r , $x_{\text{same.field}}(s, r)$ is an indicator of whether s and r work in the same field, $x_{\text{messages.send}}(r, s, t)$ is the total number of messages sent by r to s until time t . When the parameters have been estimated, $\beta_{\text{hierarchical.diff}}$, $\beta_{\text{same.field}}$, and $\beta_{\text{reciprocity}}$ quantify the relative importance each of these hypothesized drivers of email interaction. For example, a positive value for $\beta_{\text{hierarchical.diff}}$ implies that employees tend to share information with colleagues of higher ranks while a negative value for $\beta_{\text{hierarchical.diff}}$ suggests a tendency for employees to share information of lower ranks.

Contrary to most studies of network dynamics ([22, 37-41]), we do not assume that the effects that drive emailing rates remain constant over time. For example, reciprocity might drive email sending throughout the day (when someone receives a request for help, he might feel obliged to respond promptly), but this may not hold during the night or weekends (when people cannot be expected to read and respond to their work emails quickly). Similarly, at the onset of new projects, the rate of messages from managers to subordinates may be high (in order to make sure the project gets started properly), but once the project is underway email may start to be sent at higher rates from subordinates to their managers (reporting their progress) while sending rates of managers to subordinates are likely to drop. At the global level, this effectively shifts activity to other places in the network (once the project has started) and reverses the directionality of the interaction. If we were to assume stationarity of the dynamics, such shifts would remain undetected, and an important characteristic of the network dynamic would be overlooked.

Therefore, rather than fitting parameters as constant across the entire event stream, we propose a moving window technique as follows:

1. Specify a 'window' of a certain length. Ideally, the length of the window should be chosen depending on the temporal nature of the effects. For example, if one is interested in the dynamic behavior over months, the window length can be set to one or two months.
2. Fit the model for the subset of relational events that took place in the first period of the window length.
3. Move the window with a small increment (such that it partly overlaps with the previous window) and fit the model to this next subset of relational events.
4. Repeat step 3 until the last period of relational events has been analyzed.

Note that sufficient numbers of events should occur within any given window to properly fit the statistical model. When the window is chosen too small with too few events observed within a window, the estimation of the network effects can become unstable. Therefore, window length should depend on the time scale and sample size.

In Figure 1 the moving window technique in the relational event model is illustrated to get an idea of how the popularity of newcomers as receivers of messages changes over time. Initially, “old-timers” are the most popular receivers but as time goes by this gradually shifts to newcomers. This shift can be explained by the fact that it often takes some time before newcomers integrate into an organization. By investigating how the popularity of newcomers changes over time, using the proposed moving window approach in the relational event model, we get a better understanding of how and how fast this integration process occurs in an organization.

[Figure 1 about here]

3. Email network: empirical analysis illustration, part I

We illustrate the first part of our proposed methodology by applying the relational event model to a relational event history of email messages collected in a large consultancy firm. In particular, these were email messages about innovation projects and ideas about innovation activities. The company extracted the email messages from the entire body of email messages among its employees, by automated scanning of the text of the subject header and body of each email message and categorizing whether it was about “innovation.” The messages were then anonymized, and we obtained the event stream of these innovation-related email messages. Also, we received some high-level information about the individuals, such as location, department, expertise, and tenure. In addition, we were provided with some information about events in the organization and the undertaking of several innovation-related projects. We use this information in our interpretation of some of the effects below.

The firm was particularly interested how its employees communicated (if at all) about innovation because the firm wanted to stimulate innovation in the firm and this was a topic that was quite low on the minds of the firm’s employees. Hence, it was of great relevance to the organization to know where in the firm’s network innovation tended to be on the agenda and how they could understand where and when its employees would discuss innovation. The relational event network that is analyzed here consisted of the 70 most active employees sending 2081 innovation-related messages among each other over the course of a year. Idle periods, such as midnight hours, were discarded from the data. Table 1 displays the first six relational events. Figure 2 contains several descriptive statistics of the data: the tenure of the employees, the distribution of the employees over the six divisions, the hierarchical position of employees on a 1 (e.g., “secretary”) to 4 (e.g., “firm partner”) scale, the employee distribution over their six different building locations, the distribution of the hierarchical difference, and the number of messages sent each month. To get an initial idea about the interaction structure between the employees and how this changes over the months in the year, Figure 3 shows aggregated networks where the thickness of the lines is proportional to the number of messages sent between employees. Nodes have different shapes, sizes, filled colors, and colored lines to reflect the different buildings the employees work at, their hierarchical positions, the divisions they work at, and whether they have less than four years of tenure or not, respectively. The figure gives an initial idea of how the communication structures about innovation-related topics gradually change over the year. It is common to depict longitudinal networks by a series of plots of subsequent time periods. The figure shows that the communication is dynamic, but it does not make clear exactly

how the network changes, how fast it changes, and what the drivers are that explain why some parts of the network are more dynamic than others. Especially when the network data consist of streams of continuous events in real-time, collapsing data into a series of (arbitrary) intervals does not uncover the real dynamics underlying the network. This is also one of the reasons why the relational event approach is better fit for continuous time event data in networks than alternative approaches to modeling networks over time, such as the stochastic actor-oriented approach of Siena [16, 17, 42] or temporal versions (TERGM, STERGM) of the ERGM model [43]. The Siena and (S)TERGM models consider a series of observations of the network, such as the 12 we have plotted here, and models how the network changes from one month to the next. Since we are dealing with networked events that flow continuously, timed by the second, these models are not applicable to this type of data and event-based (rather than state-based) network models are required.

[Figures 2 and 3 about here]

In the relational event model we fit here, we model the rate at which vertex s sends an innovative email to vertex r at time t as a log-linear function of:

- the hierarchical difference between sender and receiver (calculated as the hierarchical level of the receiver minus the level of the sender);
- whether the sender and receiver work in the same building (1 = yes, 0 = no);
- whether the sender and receiver work in the same division (1 = yes, 0 = no);
- the hierarchical level of the sender (multiple levels on a linear scale);
- whether the sender is considered an “old-timer” (i.e., someone who has worked at least four years at the company; 1 = yes, 0 = no);
- whether receiver is considered a “old-timer” (0 = yes, 1 = no);
- how recently the sender sent his last message (days)
- how many messages the sender received from the receiver in the past (number);
- how many messages the sender has received from others than the current recipient (this represents someone’s bridging behavior; number).

Table 2 includes the full list of the statistics and their definitions which were based on [4,22,35]. Maximum likelihood estimates and standard errors were obtained using the ‘coxph’-function in the R-package ‘survival’ [44, 45].

[Tables 1 and 2 about here]

When predicting the future by taking into account the past, it is important to consider how long the relevant past is likely to be. When actors in the network operate on short memories and tend to only respond to fairly recent events, the history to be included in the model should be equally short. In real-life cases, it is rarely known for how long a sent email message might affect future emailing rates. Therefore, we fitted the moving model separately for a history length of 60 days, to investigate short-term dynamics, and 150 days, to investigate long-term dynamics. In the shorter history, the current emailing rate from s to r only takes into account the emails that were sent and received in the network within the past 60 days. When operating under longer history, email exchanges of the past 150 days (approximately 5 months) matter. This allows us to analyze how the dynamics are affected by short-term or longer-term histories. Of course, other lengths could be considered as well (e.g., see [46]). Unfortunately, there is virtually no literature to inform a theory-based choice for which history-lengths to consider and any choice would be, to some extent, arbitrary. Since we have a year of data, memory lengths of much more than 150 days (which is almost half a year) would not be useful. The shorter the period considered, the lower number of events that would occur in it, which affects the

feasibility of fitting a model that contains a large set of variables. Our 60-day interval generated windows with sufficient numbers of events for the models to be stable.

[Figures 4 and 5 about here]

Figure 4 shows the findings from the model, for the 60-days memory. On the horizontal axis, the months along the year are shown. The first time point contains the days of January and February, since this is the first point during the year that a 60-day history is available. Figure 5 shows the same results for a 150-day memory. The top row of Figure 4 shows dyadic characteristics of the sender → receiver dyad. Top-left shows the effect of the hierarchical difference between sender and receiver. Overall, the effect is positive throughout the year: rates of sending from s to r increase as the receiver has a higher hierarchical level than the sender. In other words: email messages regarding innovation tend to be sent up the hierarchy at higher rates than messages sent down the hierarchy. This effect increases at the beginning of the year and slightly decreases during the Summer period when the rates of emails are less dominated by hierarchically upward travel. An important issue in many large consulting firms is how to distribute the consultants across the consulting firm's offices. Communication theory suggests that communication between people can be strongly driven by physical proximity between them, especially for innovation-related communication [47-51]. However, consultants perform much of their work on the premises of external clients. Hence, it can be argued that the communication patterns of the consultants should not be affected so much by the location of their official office location. In line with this, the variable "same building" shows a positive effect around the summer months (when these consultants tend to be more in their own offices) and is negligible in the beginning and end of the year (showing that the consultants no longer favor communication with same-building colleagues). Simultaneously, employees of the firm tend to send emails at considerably higher rates to colleagues of the same division, compared to cross-division communication. In other words, across the year expertise (cognitive) proximity is a stronger and more consistent driver of email events than geographic proximity. Overall, there is some variation across the year, and the dynamics tend to change in the summer months and at the beginning and end of the year. The rate at which messages are sent to a receiver in the same building is, overall, only about 10 percent higher than the rate of message-sending to receivers in another building; being in the same division on average multiplies the email rate approximately fivefold. These effects are additive, so a dyad in which the receiver has a higher hierarchical level than the sender and both belong to the same division has a higher rate than a same-division same-level dyad. A relational event network approach can easily accommodate such heterogeneous communication rates.

The middle rows of Figures 4 and 5 show the effects of the actor-specific characteristics (rather than dyad-level characteristics). The left figures show that rates tend to be lower as the hierarchical position of the sender increases, during the spring and fall. In the summer and winter of this year, hierarchy did not make a difference in sending rate. This is an interesting finding, as it shows how the organization responded to new innovative projects. While there are always new projects starting or old project ending at any point during the year somewhere in the organization, two innovation projects started that addressed the organization as a whole: one started at the beginning of May, another in the second half of November. Even though these projects were a lot smaller than most of the typical projects that occurred throughout the entire year across the organization, they triggered an atypical interaction pattern. Whereas innovation-related emails tended to be sent mostly by lower-hierarchy employees during the "regular" parts of the year, the start of these small innovation projects triggered high rates of interaction by those higher up the hierarchy; this makes higher-hierarchy employees gradually more dominant in the interaction. Perhaps even more interesting is that these

projects took many months, but the higher-hierarchy-employee dominance only lasted for a while; the interaction network reverted to favoring lower-hierarchy-employee as senders of innovation-related messages before the innovation project was even halfway. When only a short memory is considered (Figure 4), the response of the interaction patterns to these small-scale interventions show clearly. When longer histories are used to explain the rates, the results emphasize the routine of lower-hierarchy employees dominating email sending throughout the year.

An interesting dynamic in the firm considers the role of tenure. We found clear indications of different roles between the so-called “newcomers” (those who had been in the firm less than four years at the beginning of the year) compared to the “old-timers” (those with tenures of four years or more). The middle and right Figures 4 and 5 show these effects. Those newcomers showed higher rates of sending innovation-related emails than the old-timers. This is in line with the suggestion in the literature that increased tenure tends to make employees less innovative and more prone to maintain their routines. During the year, the relative prominence of newcomers over old-timers gradually decreases, which is partly due to newcomers gaining tenure and slowing down their innovative pace. We find a related effect concerning whom is on the receiving end of the innovation-related email. At the beginning of the year, the tenure effect is not (short-term) or barely (long-term) statistically significant, while at the end of the year newcomers start to receive email at higher rates than the old-timers. The effects are not strong, but do show a consistent story across the observation year. Together these results show that newer employees are more active in discussing innovation than older employees and this effect diminishes as employees become more established in the firm.

Interestingly, although newer employees may be more suited to think outside-the-box and be active in communicating about it, they are not considered particularly worthy receivers of such interaction. As their tenure increases during the year, the newer employees tend to receive more innovation messages. In addition to increased tenure, it is also possible that new (innovative) projects increasingly involved younger employees throughout the year, but we do not have the data to test this. Overall, the analysis shows that hierarchy (here regarding organizational tenure) plays a role in explaining and modeling heterogeneity in interaction rates over time in a network.

The bottom row of Figures 4 and 5 focuses on drivers of email intensity at the level of the network as a whole. The most common variable used in network research is that of reciprocity. The 150-day memory results indicate that there is a norm of reciprocity in the firm that makes people take into account from whom they received messages. Faster responses are given to senders the more the employee had received innovation-related messages from them in the past. Note, however, that the estimates for the short memory windows have very wide confidence bounds in the first half of the year, which is due to the lower number of events present in those short memory windows. For the long memory window, the effect is positive and statistically significant for most of the year. On average, the more messages s has received from r in the past, the faster s will respond to r next.

The “bridging” statistic captures the extent to which an employee tends to send messages to those from whom he did not recently receive messages. This reflects “handing off” behavior, where the receiver of a message next spreads the ideas to others in the organization. In both Figures 4 and 5, this effect is consistently negative (and statistically significant): messages are sent at higher rates those employees recently received messages from, rather than introducing additional partners into the conversation. *Ceteris paribus*, this type of behavior can induce the forming of cohesive subgroups where people keep favoring communication among each other, rather than being tempted to turn around and engage others into the discussion who are not part of their subgroup. This effect

is quite strong, especially for organizational members operating under short memories. As innovation was put more prominently on the agenda by the management of the organization throughout the year, this tendency reduced but did remain. This makes it harder for innovative ideas to spread quickly throughout the organization; instead, they tend to keep trapped within cohesive sets of actors.

As is common in many human systems, the actors in the email network seem to have adopted an emailing routine. The bottom-left figures show that, at time t , those who sent a message shortly before time t will sooner send the next message than those who had sent their last message longer ago. In other words, the more recent the last email sending activity, the quicker the next message is sent; the more distant their last emailing activity, the longer it will take until they send their next message. Regarding email sending, this shows that the employees tend to keep doing what they did recently. This measure of routine was quite strong and quite stable over the entire observation period.

4. Quantifying statistical evidence

Although the parameters of the statistical model were all fitted simultaneously, the discussion above interprets each statistic separately, conditional on the effects of all the others. In order to more fully understand the dynamics of the event model and/or to test theoretically interesting hypotheses, we need to go beyond the statistical significance of each statistic (and its fitted effect size); rather, we want to determine exactly how much evidence there is in the data for each statistic. The “Bayes factor,” a Bayesian statistical quantity, is well-suited for this purpose as the Bayes factor is a quantification of relative statistical evidence between competing statistical models [52, 53]. Bayes factors have several useful properties that are not shared by classical p -values [54]. Most importantly for our purpose, Bayes factors can be used to quantify relative evidence in favor of one model versus another. For example, a Bayes factor of 10 for a model M_0 against a model M_1 10 times more plausible that model M_0 is the data-generating model than model M_1 . Classical p -values, on the other hand, cannot quantify the evidence in favor of a null model; they can only be used to falsify the null model.

By design, it is straightforward to use Bayes factors to test multiple competing statistical models against one another simultaneously. These models can be quite complex and can even contain expectations regarding the ordering of the relative effects in the model [55]. This allows us to draw statistical conclusions about complex interactions of the variables that drive interaction rates in the event network.

As an example of this approach, we consider how one could test the order of the strength of the effect of various kinds of similarity on the email interaction rates. Research has shown that experienced cognitive similarity affects interpersonal communication, but it has also been shown that there are multiple bases of similarity (more commonly called “homophily” in social network research) that need not all affect communication equally (e.g., [56-60]). Above, we found that the hierarchical level of the sender affects the email-sending rate negatively and that email rates increase when sender and receiver are members of the same division and when they are located in the same building. Suppose we are interested in testing which effect is the largest (and which the second-largest). Also, we do not want to assume that this ordering is constant over time: for example, it could be that there are parts of the year in which residence in the same building matters more than division membership (e.g., if organizational activities are organized that focus on innovating the building

people work in). In our example, we formulate five competing models that we test against each other simultaneously, at each point in time throughout the observation period (using the same windows of data that we used above). Our baseline model assumes that the effects of hierarchical similarity, geographic similarity (same building), and expertise similarity (same division) are equal. We expect that innovation-related communication is most likely to occur among employees who perform similar work and have similar types of expertise; these tend to work in the same division (e.g., consulting, tax, or audit). We have less clear expectation regarding the effect of being located in the same building or occupying similar hierarchical (status) positions. Therefore, next to the baseline model, we include the following three competing models as well as a complementary model, denoted by M_c ,

$$M_0: \beta_{\text{same.division}} = \beta_{\text{same.hierarchical.level}} = \beta_{\text{same.building}}$$

$$M_1: \beta_{\text{same.division}} > \beta_{\text{same.hierarchical.level}} = \beta_{\text{same.building}}$$

$$M_2: \beta_{\text{same.division}} > \beta_{\text{same.hierarchical.level}} > \beta_{\text{same.building}}$$

$$M_3: \beta_{\text{same.division}} > \beta_{\text{same.building}} > \beta_{\text{same.hierarchical.level}}$$

M_c : none of the above models.

The complementary model covers all other possible combinations of constraints on these effects, and therefore it serves as a check to see whether the presumed models $M_0, \dots, M_3$ receive any fair amount of evidence from the data. Also, if there is considerable evidence for M_c , this would suggest that there are relevant models that are not included in $M_0, \dots, M_3$ that should be considered as well.

Whereas the interpretation of the Bayes factors themselves is insightful in itself, Bayes factors can also be translated to *posterior model probabilities* ([61]). Posterior model probabilities quantify how likely it is for each model to have generated the data given all models under consideration. These probabilities give a direct answer to the research question of which model is most likely to be true given the data and to what degree. Because of this intuitive interpretation, below we will report posterior model probabilities instead of Bayes factors themselves (the conclusions based on both methods would be equivalent). In particular, we calculate the posterior model probabilities based on the “default Bayes factor” approach of [62]. This approach requires only the maximum likelihood estimates and estimated error covariance matrices. Using this technique with the moving window we can see how statistical evidence between hypotheses or theories changes in real time.

5. Email network: empirical analysis illustration, part II

We first illustrate the Bayes factor approach by focusing on the effect of being located in the same building. The empirical results from Section 3 show that co-location does not seem to have a strong effect on interacting with innovation-related matters. We also found that the effects are dynamic over time and that co-location mattered mainly during the summer period. The Bayes factor approach enables us to investigate exactly *when* the effect of being a newcomer switches between neutral (i.e., it has no effect), positive (being in the same building increases emailing rates), and negative (being residents of the same building lowers emailing rates). We do this by formulating the three respective models:

$$M_0: \beta = 0$$

$$M_1: \beta < 0$$

$$M_2: \beta > 0,$$

with β representing the effect of sender and receiver being located in the same building. The results are shown in Figure 6 where the y-axes contain the posterior model probabilities of each model, for both memory lengths.

For the shorter memory length between February and April, there is approximately a .85 posterior probability that there is no effect of being in the same building. As the Summer months come closer, the probabilities of “no effect” and “positive effect” move closer to each other and around the middle of May the “positive effect” becomes a more likely driver of email rates than “no effect.” Throughout, the probability of the hypothesis that being in the same building decreases email exchange being correct remains very close to zero, except for January and December. After the Summer, from September on, the probability that being in the same building is the correct hypothesis drops sharply and the evidence for the three hypotheses gradually becomes similar as in the beginning of the year. The longer memory lengths show the same dynamic, but a bit more pronounced.

These figures are useful to show how the direction of effects alternate, to recognize criticality in the event network [31], to model when and why switches of the direction of effects occur. For example, we can get a better understanding of how the integration process in organizations occurs by investigating the network effects of employees who are dispersed or co-located.

[Figures 6 and 7 about here]

Next, we extend the analysis by computing the statistical evidence of the five models formulated in Section 4 about the relative importance of sender and receiver similarity in location, expertise, and hierarchical level. This analysis allows a researcher to draw conclusions regarding the evolution of the relative likelihood of multiple effects simultaneously. Figure 7 shows the posterior model probabilities. Two models strongly outperform the other three consistently: models 1 and 2. Both of these models state that the effect of being in the same division is larger than the other two effects. The sum of the probabilities of these two models tends to be over 0.8 across the entire year, approaching 1.0 at many time points. When two employees are assigned to the same division, their innovation communication rate is more strongly driven by their shared affiliation than by being geographically proximate or occupying the same hierarchical level. It is also clear that models 1 and 2 compete for prominence. The positive effect of hierarchical similarity does not always outweigh that of similarity in geographic location. When a 60-day memory is assumed, the models alternate in importance, with a clear winner at each point in time—just not a clear winner throughout. For the longer memory analysis, Figure 7 shows that there are two points in time (around mid-April and mid-June) where the models come close, and the system can go either way. At the first time point, around mid-April, the effect of hierarchical similarity starts to quickly outweigh similarity in location, giving model 2 a higher probability than model 1. Another point of criticality occurs around mid-June after which status hierarchy continues its prominence over location. It is important to note that the complement model M_c consistently has a posterior model probability of almost zero. This probability is the posterior probability for all of the possible models together that are not part of $M_0, \dots, M_3$. This means that there is no further model that needs to be considered, besides the ones that are already taken into account.

Overall, regarding similarity, the results are resounding in the dominance of expertise similarity and show that geographic similarity, despite the many studies in the literature (e.g., [59, 63-67])

suggesting a strong positive effect, does matter, but only in third place after similarity in expertise and hierarchy.

6. Concluding remarks

This paper shows how time-sensitive social network interaction streams can be analyzed using a dynamic relational event model with a moving window technique. By setting the window length to a short period we can zoom in on the drivers of the interaction process when the network members are assumed to only respond to recent interaction history and by setting the window length to a larger period we can see the effects of them operating under longer memory of past interaction [4]. A Bayes factor procedure was proposed to investigate how statistical evidence between multiple theories evolves. This procedure can aid the development of time-sensitive theory on social network dynamics, which are currently underdeveloped in the literature [4-7, 68-70].

The methodology was illustrated using an event history of email messages between colleagues in a large consultancy firm. The analysis showed how exogenous drivers, such as whether sender and receiver work in the same division, have similar hierarchical levels, or work in the same location, affect the interaction process. We observed that the importance of these drivers could change substantially over time. The methodology can uncover points of criticality of the interaction system [31] and shows which variables play a role in this. Also, findings of this type of study are vital in the effective design of interaction networks. For example, managers of innovative teams who want to stimulate knowledge-sharing might be less concerned with putting project members in the same building but should be more concerned with the effects of hierarchy and cognitive and task similarity. More studies are needed to draw more definite conclusions, but findings of the real-time development and drivers of communication in systems of real people are very scarce and are vital to inform future time-sensitive theory.

For future research, it would be useful to extend the model by assuming dynamic network drivers using auto-regressive models or state-space models. This allows us to directly model the dynamic nature of the network drivers in real time. Such an approach is also likely to result in more stable estimates, particularly for short-term effects. Depending on the scale of the dynamic behavior, the effects can be modeled to change either every day, every week, every month or over longer periods. To fit such a dynamic model, classical methods ([71]) or Bayesian methods ([72]) could be considered. The Bayesian approach may be particularly suitable due to the natural updating of the dynamic model when observing new observations using Bayes' theorem. The challenge of this extension would be to develop efficient estimation algorithms for fitting and updating the model. In a Bayesian framework, Markov chain Monte Carlo (MCMC) techniques may be too slow for this purpose. Alternatively, sequential Monte Carlo techniques ([73]) may be a promising approach for these dynamic relational event models.

FIGURE 1: GRAPHICAL REPRESENTATION OF THE MOVING WINDOW

As an example, the graph shows how the popularity of newcomers as receivers changes over time. Events are shown as a “|” along the horizontal time axis. Above the axis, three instances of the moving window are shown. The graph at the bottom shows the parameter estimate over time: in the beginning newcomers are less popular receivers (negative effect) but over time newcomers become more popular receivers (the effect turns positive).

FIGURE 2: DESCRIPTIVE STATISTICS OF THE EMPIRICAL DATA. REGARDING THE NUMBER OF YEARS EMPLOYED, 50% OF THE 70 EMPLOYEES HAD 4+ YEARS OF TENURE (“OLDTIMERS”).

FIGURE 3: AGGREGATED NETWORKS OVER DIFFERENT MONTHS. AN OLDTIMER (NEWCOMER) IS AN EMPLOYEE HAVING 4+ (LESS THAN 4) YEARS OF TENURE.

FIGURE 4. ESTIMATED NETWORK EFFECTS OVER TIME USING A MOVING WINDOW OF 60 DAYS

The dashed lines represent 95%-confidence bounds.

FIGURE 5. ESTIMATED NETWORK EFFECTS OVER TIME USING A MOVING WINDOW OF 150 DAYS.

The dashed lines represent 95%-confidence bounds.

FIGURE 6. TRENDS OF POSTERIOR PROBABILITIES OF WORKING IN THE SAME BUILDING

Trends of posterior probabilities of no effect (solid line) versus a negative effect (dashed line) versus a positive effect (dotted line) of residing in the same building on receiving for a moving window of 60 days (upper panel) and 150 days (lower panel).

FIGURE 7. TRENDS OF POSTERIOR PROBABILITIES BETWEEN MULTIPLE COMPETING MODELS

The lines represent the posterior probabilities of statistical models assuming different orderings of the importance of similarity in division, hierarchy, and building of sender and receiver, for a moving window of 60 days (left panel) and 150 days (right panel).

TABLE 1. FIRST SIX RELATIONAL EVENTS OF THE CONSULTANCY EMAIL DATA. THE DATES ARE FORMATTED AS MM/DD/YYYY.

Sender	Receiver	date	time
1	7	01/06/2010	07:13:13
1	2	01/06/2010	07:13:53
1	7	01/06/2010	09:10:02
1	2	01/06/2010	09:11:32
8	12	01/07/2010	13:42:21
8	11	01/07/2010	13:43:71

TABLE 2. OVERVIEW OF STATISTICS AND CORRESPONDING EFFECTS IN THE RELATIONAL EVENT MODEL OF THE EMPIRICAL APPLICATION.

Statistic	Effect
Hierarchical level of the sender is measured on a scale of 1 to 4 (i.e., a secretary has level 1, and a partner has level 4).	$\beta_{hierarchy.sender}$: A positive (negative) effect implies that employees with a high hierarchical level are more (less) active senders.
Hierarchical difference between sender s and receiver r is the hierarchical level of r minus the hierarchical level of s .	$\beta_{hierarchical.diff}$: Positive (negative) effect implies that sending rates increase (decrease) as receivers have a higher hierarchical level.
Same building indicates whether the sender and receiver work in the same building (1=same building; 0=different buildings).	$\beta_{same.building}$: A positive (negative) effect implies sending rates increase (decrease) as receivers reside in the same building as the sender. A zero effect means that being in the same building does not affect emailing rates.
Same division indicates whether the sender and receiver work in the same division (e.g., consultancy, tax, audit) (1=same division; 0=different divisions).	$\beta_{same.division}$: A positive (negative) effect implies sending rates increase (decrease) as receivers are a member of the same division as the sender. A zero effect means that membership of the same division does not affect emailing rates.
Sender with 4+ years tenure indicates whether the sender has worked at least four years at the company (at the beginning of the year).	$\beta_{sender.middle.tenure}$: A positive (negative) effect implies that employees who have worked at least 4 years at the company have higher (lower) message-sending rates.
Receiver with max three years tenure indicates whether the receiver has worked less than four years at the company (at the beginning of the year)..	$\beta_{receiver.begin.tenure}$: A positive (negative) effect implies that employees who have worked between 0 and 3 years at the company are (un)popular receivers.
Recency is a dyadic statistic quantifying how recent a sender s sent a message in the past. It is computed as $\frac{1}{t(s)+1}$ [22].	$\beta_{recency.send}$: A positive (negative) effect implies that sending rates increase (decrease) the more recent an employee had sent their last message.
Reciprocity is a dyadic statistic quantifying how many messages a receiver r sent to the sender s in the past. It is computed as $\log\left(\frac{\#messages(r,s,t)+1}{N_t+n(n-1)}\right)$, where N_t is the number of	$\beta_{reciprocity}$: A positive (negative) effect implies that senders have higher (lower) rates of sending messages to receivers from whom they received more messages in the past.

messages sent until time t , and n is the number of nodes [22].	
<p>Bridging is a dyadic statistic quantifying how many messages the sender received from other nodes than the receiver r in the past. It is computed as $\log\left(\frac{\#messages(not(r),s,t)+1}{N_t+n(n-1)}\right)$, where N_t is the number of messages sent until time t, and n is the number of nodes [22].</p>	<p>$\beta_{bridging}$: A positive effect implies that senders have a higher rate of redirecting a conversation to employees who did not send messages to the sender in the past. A negative effect implies higher rates of sending to those the sender received past messages from.</p>

REFERENCES

- [1] Snijders TAB, Baerveldt C. A multilevel network study of the effects of delinquent behavior on friendship evolution. *The Journal of Mathematical Sociology*. 2003;27:123-51.
- [2] Stadtfeld C, Hollway J, Block P. Dynamic Network Actor Models: Investigating Coordination Ties through Time. *Sociological Methodology*. 2017;47:1-40.
- [3] Leenders RTAJ. *Structure and Influence: Statistical Models for the Dynamics of Actor Attributes, Network Structure, and Their Interdependence*. Amsterdam: Tesla Thesis Publishers; 1995.
- [4] Leenders RTAJ, Contractor NS, DeChurch LA. Once upon a time: Understanding team processes as relational event networks. *Organizational Psychology Review*. 2016;6:92-115.
- [5] Kozlowski SWJ, Chao GT. The Dynamics of Emergence: Cognition and Cohesion in Work Teams. *Managerial and Decision Economics*. 2012;33:335-54.
- [6] Kozlowski SWJ, Chao GT, Grand JA, Braun MT, Kuljanin G. Advancing Multilevel Research Design: Capturing the Dynamics of Emergence. *Organizational Research Methods*. 2013;16:581-615.
- [7] Kozlowski SWJ, Chao GT, Grand JA, Braun MT, Kuljanin G. Capturing the multilevel dynamics of emergence: Computational modeling, simulation, and virtual experimentation. *Organizational Psychology Review*. 2016;6:3-33.
- [8] Snijders TAB. Methods for longitudinal social network data: review and markov process models. *MultivarStatistics*. 1995:211-27.
- [9] Cronin MA. Advancing the science of dynamics in groups and teams. *Organizational Psychology Review*. 2015;5:267-9.
- [10] Cronin MA, Weingart LR, Todorova G. Dynamics in Groups: Are We There Yet? *Academy of Management Annals*. 2011;5:571-612.
- [11] Brandenberger L. Trading favors—Examining the temporal dynamics of reciprocity in congressional collaborations using relational event models. *Social Networks*. 2018;54:238-53.
- [12] Gross T, Blasius B. Adaptive coevolutionary networks: a review. *Journal of The Royal Society Interface*. 2008;5:259-71.
- [13] Gorochoowski TE, Bernardo MD, Grierson CS. Evolving dynamical networks: A formalism for describing complex systems. *Complexity*. 2012;17:18-25.
- [14] Zečević AI, Šiljak DD. Dynamic graphs and continuous Boolean networks, I: A hybrid model for gene regulation. *Nonlinear Analysis: Hybrid Systems*. 2010;4:142-53.
- [15] Boccaletti S, Latora V, Moreno Y, Chavez M, Hwang DU. Complex networks: Structure and dynamics. *Physics Reports*. 2006;424:175-308.
- [16] Snijders TA, Van de Bunt GG, Steglich CE. Introduction to stochastic actor-based models for network dynamics. *Social Networks*. 2010;32:44-60.
- [17] Snijders TAB. Siena: Statistical Modeling of Longitudinal Network Data. In: Alhajj R, Rokne J, editors. *Encyclopedia of Social Network Analysis and Mining*. New York, NY: Springer New York; 2014. p. 1718-25.
- [18] Caimo A, Friel N. Actor-Based Models for Longitudinal Networks. In: Alhajj R, Rokne J, editors. *Encyclopedia of Social Network Analysis and Mining*. New York, NY: Springer New York; 2014. p. 9-18.
- [19] Krivitsky PN, Handcock MS. A separable model for dynamic networks. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2014;76:29-46.
- [20] Hanneke S, Fu W, Xing EP. Discrete temporal models of social networks. *Electron J Statist*. 2010;4:585-605.
- [21] Blonder B, Dornhaus A. Time-Ordered Networks Reveal Limitations to Information Flow in Ant Colonies. *PLoS ONE*. 2011;6:e20298.
- [22] Butts CT. A Relational Event Framework for Social Action. *Sociological Methodology*. 2008;38:155-200.
- [23] McGrath JE, Arrow H, Berdahl JL. The Study of Groups: Past, Present, and Future. *Personality and Social Psychology Review*. 2000;4:95-105.

- [24] Gabbay SM, Leenders RTAJ. CSC: The Structure of Advantage and Disadvantage. In: Leenders RTAJ, Gabbay SM, editors. *Corporate Social Capital and Liability*. New York: Wolters-Kluwer Academic Publishers; 1999. p. 1-14.
- [25] McGrath JE, Argote L. Group processes in organizational contexts. In: Hogg MA, Tindale RS, editors. *Blackwell Handbook of Social Psychology: Group Process*. Oxford, UK: Blackwell; 2001. p. 603-27.
- [26] Brenner MH. Small networks, evolution of knowledge and species longevity: Theoretical integration and empirical test. *Chaos, Solitons & Fractals*. 2017;104:314-22.
- [27] Liu Y-Y, Slotine J-J, Barabási A-L. Controllability of complex networks. *Nature*. 2011;473:167.
- [28] Gorochoowski TE, Grierson CS, di Bernardo M. Organization of feed-forward loop motifs reveals architectural principles in natural and engineered networks. 2018;4.
- [29] Richardson TO, Gorochoowski TE. Beyond contact-based transmission networks: the role of spatial coincidence. *Journal of The Royal Society Interface*. 2015;12:20150705.
- [30] Masuda N, Holme P. *Temporal Network Epidemiology*. 1 ed. Singapor: Springer Singapore; 2017.
- [31] Grigolini P. Call for papers: Special issue on evolutionary game theory of small groups and their larger societies. *Chaos, Solitons & Fractals*. 2017;103:371-3.
- [32] Brandes U, Lerner J, Snijders TAB. *Networks Evolving Step by Step: Statistical Analysis of Dyadic Event Data*. *ASONAM '09 Proceedings of the 2009 International Conference on Advances in Social Network Analysis and Mining (20-22 July 2009, Athens, Greece)*: IEEE Computer Society; 2009. p. 200-5.
- [33] Cox DR. Regression Models and Life-Tables. *Journal of the Royal Statistical Society Series B (Methodological)*. 1972;34:187-220.
- [34] Box-Steffensmeier JM, Jones BS. Time is of the Essence: Event History Models in Political Science. *American Journal of Political Science*. 1997;41:1414-61.
- [35] DuBois C, Butts CT, McFarland D, Smyth P. Hierarchical models for relational event sequences. *Journal of Mathematical Psychology*. 2013;57:297-309.
- [36] Lawless JF. *Statistical Models and Methods for Lifetime Data*. 2nd ed: Wiley; 2003.
- [37] Quintane E, Conaldi G, Tonellato M, Lomi A. Modeling Relational Events: A Case Study on an Open Source Software Project. *Organizational Research Methods*. 2014;17:23-50.
- [38] Pilny A, Schecter A, Poole MS, Contractor N. An illustration of the relational event model to analyze group interaction processes. *Group Dynamics: Theory, Research, and Practice*. 2016;20:181-95.
- [39] Schecter A, Pilny A, Leung A, Poole MS, Contractor N. Step by step: Capturing the dynamics of work team process through relational event sequences. *Journal of Organizational Behavior*. 2017:n/a-n/a.
- [40] Stadtfeld C, Geyer-Schulz A. Analyzing event stream dynamics in two-mode networks: An exploratory analysis of private communication in a question and answer community. *Social Networks*. 2011;33:258-72.
- [41] Tranmer M, Marcum CS, Morton FB, Croft DP, de Kort SR. Using the relational event model (REM) to investigate the temporal dynamics of animal social networks. *Animal Behaviour*. 2015;101:99-105.
- [42] Steglich C, Snijders TAB, West P. Applying SIENA: An Illustrative Analysis of the Coevolution of Adolescents' Friendship Networks, Taste in Music, and Alcohol Consumption. *Methodology: European Journal of Research Methods for the Behavioral and Social Sciences*. 2006;2:48-56.
- [43] Block P, Koskinen J, Hollway J, Steglich C, Stadtfeld C. Change we can believe in: Comparing longitudinal network models on consistency, interpretability and predictive power. *Social Networks*. 2018;52:180-91.
- [44] Therneau TM. *A Package for Survival Analysis in S*. 2.38 ed 2015.
- [45] Therneau TM, Grambsch PM. *Modeling Survival Data: Extending the {C}ox Model*. New York: Springer; 2000.
- [46] Quintane E, Pattison PE, Robins GL, Mol JM. Short- and long-term stability in organizational networks: Temporal structures of project teams. *Social Networks*. 2013;35:528-40.
- [47] Allen TJ. Communication Networks in R&D Labs. *R&D Management*. 1971;1:14-21.
- [48] Zander U, Kogut B. Knowledge and the Speed of the Transfer and Imitation of Organizational Capabilities: An Empirical Test. *Organization Science*. 1995;6:76-92.
- [49] Hoegl M, Proserpio L. Team member proximity and teamwork in innovative projects. *Research Policy*. 2004;33:1153-65.

- [50] Chong DSF, Eerde W, Rutte CG, Chai KH. Bringing Employees Closer: The Effect of Proximity on Communication When Teams Function under Time Pressure. *Journal of Product Management*. 2012;29:205-15.
- [51] Kraut RE, Fussell SR, Brennan SE, Siegel J. Understanding effects of proximity on collaboration: Implications for technologies to support remote collaborative work. In: Hinds P, Kiesler S, editors. *Distributed work*. Cambridge, MA, US: MIT Press; 2002. p. 137-62.
- [52] Kass RE, Raftery AE. Bayes Factors. *Journal of the American Statistical Association*. 1995;90:773-95.
- [53] Jeffreys H. *Theory of Probability*. 3 ed. Oxford: Oxford University Press; 1961.
- [54] Wagenmakers E-J. A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*. 2007;14:779-804.
- [55] Mulder J. Bayes factors for testing inequality constrained hypotheses: Issues with prior specification. *British Journal of Mathematical and Statistical Psychology*. 2014;67:153-71.
- [56] Festinger L. A theory of social comparison processes. *Human relations*. 1954;7:117-40.
- [57] Monge PR, Contractor NS. *Theories of Communication Networks*. New York: Oxford University Press; 2003.
- [58] VandenBulte C, Moenaert RK. The Effects of R&D Team Co-location on Communication Patterns among R&D, Marketing, and. *Management Science*. 1998;44:S1-S18.
- [59] Huang Y, Shen C, Contractor NS. Distance matters: Exploring proximity and homophily in virtual world networks. *Decision Support Systems*. 2013;55:969-77.
- [60] McPherson M, Smith-Lovin L, Cook JM. Birds of a Feather: Homophily in Social Networks. *Annual review of sociology*. 2001;27:415-44.
- [61] Braeken J, Mulder J, Wood S. Relative effects at work: Bayes factors for order hypotheses. *Journal of Management*. 2015;41:544-73.
- [62] Gu X, Mulder J, Hoijtink H. Approximated adjusted fractional Bayes factors: A general method for testing informative hypotheses. *British Journal of Mathematical and Statistical Psychology*. 2018;71:229-61.
- [63] Allen TJ. *Managing the Flow of Technology*. Boston: MIT Press; 1977.
- [64] Jones BF, Wuchty S, Uzzi B. Multi-University Research Teams: Shifting Impact, Geography, and Stratification in Science. *science*. 2008;322:1259-62.
- [65] Song M, Berends H, Van Der Bij H, Weggeman M. The Effect of IT and Co-location on Knowledge Dissemination. *Journal of Product Innovation Management*. 2007;24:52-68.
- [66] Olson GM, Olson JS. Distance Matters. *Human-Computer Interaction*. 2000;15:139-78.
- [67] Kabo F, Hwang Y, Levenstein M, Owen-Smith J. Shared Paths to the Lab: A Sociospatial Network Analysis of Collaboration. *Environment and Behavior*. 2015;47:57-84.
- [68] Mitchell TR, James LR. Building Better Theory: Time And The Specification Of When Things Happen. *Academy of Management Review*. 2001;26:530-47.
- [69] Monge PR. Theoretical and Analytical Issues in Studying Organizational Processes. *Organization Science*. 1990;1:406-30.
- [70] Kozlowski SWJ. Advancing research on team process dynamics: Theoretical, methodological, and measurement considerations. *Organizational Psychology Review*. 2015;5:270-99.
- [71] Hamilton JD. *Time Series Analysis*. Princeton, NJ: Princeton University Press; 1994.
- [72] West M, Harrison J. *Bayesian forecasting and dynamic models*. New York: Springer-Verlag; 1997.
- [73] Paige B, Wood F, Doucet A, Teh YW. Asynchronous anytime sequential monte carlo. *Advances in Neural Information Processing Systems* 2014. p. 3410-8.