This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Backlund, Ville-Pekka; Saramäki, Jari; Pan, Raj Kumar

Effects of temporal correlations on cascades: Threshold models on temporal networks

Published in:
Physical Review E

DOI:
[10.1103/PhysRevE.89.062815](https://doi.org/10.1103/PhysRevE.89.062815)

Published: 01/01/2014

Document Version
Publisher's PDF, also known as Version of record

Please cite the original version:
Backlund, V-P., Saramäki, J., & Pan, R. K. (2014). Effects of temporal correlations on cascades: Threshold models on temporal networks. *Physical Review E*, 89(6), 1-8. [062815].
<https://doi.org/10.1103/PhysRevE.89.062815>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Effects of temporal correlations on cascades: Threshold models on temporal networksVille-Pekka Backlund,^{*} Jari Saramäki,[†] and Raj Kumar Pan[‡]*BECS, School of Science, Aalto University, P.O. Box 12200, FI-00076 AALTO, Finland*

(Received 5 March 2014; published 26 June 2014)

A person's decision to adopt an idea or product is often driven by the decisions of peers, mediated through a network of social ties. A common way of modeling adoption dynamics is to use threshold models, where a node may become an adopter given a high enough rate of contacts with adopted neighbors. We study the dynamics of threshold models that take both the network topology and the timings of contacts into account, using empirical contact sequences as substrates. The models are designed such that adoption is driven by the number of contacts with different adopted neighbors within a chosen time. We find that while some networks support cascades leading to network-level adoption, some do not: the propagation of adoption depends on several factors from the frequency of contacts to burstiness and timing correlations of contact sequences. More specifically, burstiness is seen to suppress cascade sizes when compared to randomized contact timings, while timing correlations between contacts on adjacent links facilitate cascades.

DOI: [10.1103/PhysRevE.89.062815](https://doi.org/10.1103/PhysRevE.89.062815)

PACS number(s): 89.75.Hc, 02.50.Ey, 05.45.Tp

I. INTRODUCTION

We live in a highly complex and interdependent world where every day we make a number of decisions. Although individuals are free to make their decisions independently, it is a common human tendency to match one's attitude, beliefs, and behavior to those of his or her acquaintances [1–4]. Because of this, the ways how our social ties are structured and interlinked become important for understanding processes of social influence, especially system-level phenomena like the emergence of fads or adoption of ideas or behaviors. Such processes can then be approached quantitatively with models that define how individuals react to their neighbors' behavior, mediated through the links of some chosen type of social network where the individuals are embedded.

The most elementary models of social influence allow only two states for each individual (that is, network node): e.g., the node has either adopted an idea or it has not. Then, the rules obeyed by each node dictate how they respond to their neighbors' choices. At the very simplest, social influence can then be modeled so that, for a nonadopter, a single adopted neighbor can infect the node with adoption, leading to the standard epidemiological models such as SI (susceptible-infectious) or SIR (susceptible-infectious-recovered) [5] that adds a third possible state (recovered). Even though these models are useful for studying the basic adoption process, they lack certain crucial features needed in the context of social contagion. As an example, adopting a product or behavior is more likely if information is received from multiple adopted neighbors [6,7]. Thus a more realistic choice of the node model is to use a *threshold* model [8,9], where a node becomes an adopter only if a chosen fraction of its neighbors have already adopted, that is, where the level of peer pressure matters. The introduction of adopted individuals to a network of nonadopters may then eventually lead to a cascade where the majority of nodes become adopters; whether this happens

or not depends on the model parameters and the underlying network structure.

However, it has recently been realized that the above picture becomes more complete when an additional dimension is included: time. Social influence is transmitted through social ties, and, in reality, they are not always active—rather, there is a time series of contacts, from face-to-face conversations to emails. Incorporating the times of contacts into the network picture leads to the temporal networks framework [10], where the nodes are connected only by temporary events of short duration at specific times. Studies of time series of contacts, based on electronic communication records, have shown that contact timings display strong heterogeneity and burstiness, and that this burstiness has strong effects on dynamics taking place on temporal networks. In particular, burstiness often slows down spreading processes such as SI or SIR [11–13].

It is then natural to ask how cascades occur in threshold models if node behavior depends on the exact timings of its interactions. As there is now an additional degree of freedom—time—to be incorporated into the models, there are several possible approaches. The original threshold model by Watts [9] is purely topological, in the sense that the dynamic process of adoption takes place on a static network topology. In the other limit, the threshold models proposed by Karimi and Holme [14] and Takaguchi *et al.* [15] can be seen as purely temporal: what drives the adoption of a node is the number of recent contacts from adopted individuals, such that multiple contacts from the same adopted individual have the same effect as the same number of contacts from multiple adopted sources. Thus the role of the underlying network topology is diminished. However, one could argue that in a temporal threshold model that is explicitly designed to mimic cascades in social networks, the underlying network topology should be accounted for, as the effects of social reinforcement should be stronger when the signal arrives from multiple distinct acquaintances (for empirical results, see [6]). In other words, adoption should be more likely if several friends have already adopted.

To this end, in this paper we propose two threshold models—stochastic and deterministic—that could be categorized as temporal topological. In these models, adoption

^{*}ville-pekka.backlund@aalto.fi[†]jari.saramaki@aalto.fi[‡]raj.kumar.pan@aalto.fi

is driven by the number of contacts with a node's different network neighbors within a chosen time window representing node memory: a high number of contacts from different adopters makes adoption more likely. In this respect, the models are similar to the general model of contagion of Ref. [16]. However, our models take the actual timings of interactions into account and are therefore affected by any heterogeneities of the activity patterns of nodes, such as burstiness. Note that the thresholds are defined in terms of numbers of neighbors of nodes during the whole available time line, i.e., their static degrees. The rationale behind this assumption is the separation of time scales: we assume that the time scale of changes of the network (creation and deletion of nodes and links) is much slower than the time scale of the adoption process taking place on the network. This choice also allows for studying the effects of heterogeneities in contact timings in isolation.

Using empirical temporal network data sets, we then study with simulations the conditions leading to global cascades of adoption in these models. We found that for the stochastic model, long enough memory windows facilitated global cascades in only two out of four empirical networks. When a hard threshold is imposed on the fraction of adopted neighbors, global cascades appear only if this threshold is tuned so low that for most nodes a single contact with a single adopted neighbor triggers adoption. This is because high-degree nodes block adoption, unlike for typical spreading processes (SI, SIR). To pinpoint the effects of different types of temporal correlations, we compare the results with those obtained with reference models that remove such correlations from the empirical data. We find that, for our models, burstiness decreases adoption prevalence, whereas timing correlations between contacts on adjacent links clearly facilitate adoption.

II. METHODS

A. Basic framework and model definition

The set of $\mathcal{E}$ events representing the interactions occurring between N nodes within the time interval $[0, T]$ constitutes a temporal network. In this network, each event is denoted by a quadruplet $e \equiv (u, v, t, \delta t)$, where the event connecting nodes u and v begins at t and ends at $t + \delta t$. For instance, for call data, each event in the temporal network would correspond to a list of the caller, callee, starting time, and duration of the call. The corresponding static, aggregated network is obtained by ignoring temporal information and linking a pair of nodes if any event occurs between them in the observed period ($t \in [0, T]$). One can then view the events of the temporal networks as temporary activations of the links of the underlying static network: events define which links of the static network are active at any given point in time.

Our topological-temporal threshold models for node behavior, defined below, take into account both the timings of events and the aggregated network structure in terms of network neighborhoods. In the first model, there is an element of randomness in the adoption choices of nodes, whereas the second model is entirely deterministic. In both models, nodes have two possible states, susceptible and adopter, and initially all the nodes in the temporal network are set as susceptible

except for one randomly chosen seed adopter node. When susceptible nodes become adopters, they remain so for the rest of the time.

Stochastic threshold model. In the stochastic topological-temporal threshold model, the probability of a susceptible node becoming an adopter depends on the number of contacts from different adopted nodes in a given time window: more adopted neighbors in contact means higher probability. Formally, when a node is in contact with an adopter at time t , the number of adopted neighbors that it has been in contact with during the interval $[t - \tau, t]$ is counted. For directed networks, only incoming contacts, that is, contacts initiated by the neighbors of the node are counted. The ratio of the number of contacted neighbors to the number of neighbors in the aggregated network, i.e., the degree k_i , determines the probability of adoption:

$$\phi(i, t|\tau) := \frac{1}{k_i} \sum_{j \in v_i} \chi(j, t'), \quad t' \in [t - \tau, t]. \quad (1)$$

Here, v_i is the set of neighbors of node i in the aggregated static network. The indicator function χ is 1 if node j is an adopter and it has had at least one contact with node i within the interval $[t - \tau, t]$, and zero otherwise. The model comes with only one parameter—the time window size, or memory length, τ .

For a given node, the likelihood of adoption increases monotonically with the fraction of adopted neighbors, which is a reasonable assumption [6]. Because the threshold is defined as fractional, the effect of the number of adopted neighbors depends on node degree. Note that although the threshold rule does not directly count the number of contacts from the same adopted neighbor, these still contribute indirectly because the stochastic rule is activated whenever a contact occurs. Thus bursts of events affect the dynamics. For a schematic illustration of the model, please see Fig. 1.

Deterministic threshold model. We also consider a deterministic version of the above model. In this model, a node becomes an adopter if the ratio $\phi(i, t|\tau) \geq f$, where f is a predefined constant and common for all nodes. By varying the value of f one can control the ease with which adoption progresses. Because the threshold f is constant, the degree k_i of a node plays a major role: if $k_i \leq 1/f$, a single contact with an adopted node is sufficient for adoption. For nodes with $k_i > 1/f$, a larger number of contacts from separate adopters is required. In the deterministic model, burstiness of events on individual links has no effect unless the event trains of adjacent links are correlated: if a node does not adopt when in contact with an adopter, subsequent contacts with the same adopter are redundant and cannot cause adoption unless there are also contacts with other adopters.

Note the general difference between the models: given long enough times (and applying periodic temporal boundary conditions), the stochastic model would lead to a complete systemwide adoption (similar to the SI model, albeit much more slowly). However, in the deterministic case, even though we can tune the effective number of nodes that only require a single contact with an adopter for adoption, the diffusion dynamics may be blocked entirely by a small number of nodes where the condition for adoption is not fulfilled.

FIG. 1. (Color online) Schematic representation of the stochastic model of adoption dynamics. In the beginning of this example run, nodes A and C are adopters. Time runs from left to right, and each contact between a pair of nodes is denoted by a vertical line. The intervals τ illustrate the memory associated with each node. The lower panels show the adoption probabilities ϕ when a node is in contact with an adopter; U_1 , U_2 , and U_3 are random numbers associated with the event. For the event occurring between nodes A and D, only 1 out of 3 neighbors is an adopter and the ratio $1/3$ is smaller than the random number U_1 , and hence D does not adopt. Next, for the event between C and D, two of the neighbors of D who had contacts within the memory interval τ are adopters (A and C). Since the drawn random number $U_2 < 2/3$, D adopts. Finally, since B has only one neighbor (D), it also becomes an adopter when in contact in D, regardless of the value of the random number U_3 .

B. Data sets

In this study, we simulate the above threshold models on four different empirical temporal networks that are related to different types of human communication patterns. These data sets are as follows.

Call. The call data contain mobile phone call records of a European carrier, with ~ 759 million time-stamped voice call records over a period of 180 d. The durations of the calls are ignored in this study.

SMS. The text message data contains SMS records from the same European carrier and the same time period. There are ~ 243 million time-stamped records. We exclude data for Christmas and New Year's Eve since the behavior (in terms of the number of messages) on these two days is radically different from other days.

Email. The email data set [17] consists of logs of a university's mail server for a period of 83 d. Only interinstitute emails are considered and certain mass mailers are discarded.

Conference. The conference data set is a collection of ongoing face-to-face conversations of 113 participants of the ACM Hypertext 2009 conference for 2.5 d [18,19]. The participants were tracked with the help of radio badges that monitored their proximity to other attendees. The data contains time stamps and IDs of pairs of participants who are having a conversation within time windows of 20 s.

There are certain similarities and differences between these data sets. The first three represent electronic communication at distance, while in the conference data contacts require physical proximity. Another difference is that the call and conference conversation contacts require activity from both participants, e.g., a call needs to be picked up before it can

TABLE I. Statistical properties of the networks. From left to right: number of nodes N , number of events E , number of links m , average degree $\langle k \rangle$, average clustering coefficient $\langle c \rangle$, length of the observation period T , time resolution of the data Δt , and the directionality $\rightarrow$ or $\leftrightarrow$ of the links. $\langle k \rangle$ and $\langle c \rangle$ have been calculated for the aggregated networks by considering links as undirected.

	N	E	m	$\langle k \rangle$	$\langle c \rangle$	T	Δt	link
Call	5.8×10^6	620×10^6	15×10^6	5.0	0.23	180 d	1 s	$\leftrightarrow$
SMS	3.1×10^6	180×10^6	5.7×10^6	3.7	0.10	176 d	1 s	$\rightarrow$
Email	2993	2.0×10^5	21 736	14.5	0.21	82 d	1 s	$\rightarrow$
Conf.	113	20 818	2196	38.9	0.54	2.5 d	20 s	$\leftrightarrow$

be recorded, whereas in the SMS and email networks only the sender's activity matters. For similar reasons, we consider call and conference networks as undirected, whereas the SMS and email networks are directed. Further, in the call data one node can only participate in one call at any given point in time, whereas simultaneous contacts with multiple nodes are possible in the SMS, email, and conference networks.

In the call and SMS networks we only consider those links that have bidirectional events to ensure that nonsocial interactions, such as telemarketing, are ignored. Further, we only use the nodes in the largest connected component (LCC) of the aggregated network. For the email network the directionality of links was ignored while calculating the LCC. The properties of all data sets and corresponding aggregated networks are presented in Table I. The average degree and the average number of events per node are the highest for the conference network and lowest for the SMS network.

C. Reference models

In order to highlight the effects of different types of temporal correlations on the model dynamics, we employ the reference model approach [11,12,20]. In this approach, the original event sequences are randomized such that chosen temporal and topological characteristics are retained, while some temporal correlations are lost. Then, the outcomes of model runs using original and randomized reference networks are compared. We consider the following two null models.

Random time shuffle (RTS). The time stamps of all the events are randomly shuffled. This destroys all temporal correlations of events within links and of the nodes' activity patterns, e.g., burstiness, as well as correlations between the timings of events on adjacent links. The time variation of the global activity rate is preserved, such as the typical daily pattern.

Random offset (RO). All the events occurring on a given link are shifted by a time offset $\Delta t \in [0, T]$, chosen at random for each link. Periodic temporal boundary conditions are used, i.e., the time of each of the events is mapped back to the period $[0, T]$ by using a $\text{mod } T$ operation on the shifted time. As event sequences of different links are shifted by different offset the correlations between events on adjacent links, i.e., *link-link correlations*, are destroyed while temporal inhomogeneities within individual links, such as burstiness, are retained. By extension, all larger sequences of subsequent events, such as

FIG. 2. (Color online) Event sequence shuffling procedures for the two reference models. (a) In random time shuffling (RTS), the time stamps of events are randomly swapped. (b) In random offset (RO), all events on a given link are shifted by a random time Δt . Periodic boundaries are used, i.e., if the resulting times fall outside the observed window $[0, T]$ a $\text{mod } T$ operation maps them back to the period $[0, T]$.

temporal motifs [21,22], are also destroyed, as are the overall activation time lines of nodes.

For a schematic diagram of the shuffling procedures, see Fig. 2. Table II summarizes the key effects of the null models on temporal correlations.

D. Simulations

We study adoption dynamics in the two threshold models by simulating them using the original event sequences and the ensemble of sequences from reference models. The initial adoption is started from a random node and random event within a time period that spans the first 15% of the time period of the data set. The model rules are then iterated until the end of the event sequence. The initial time span is chosen because a large enough fraction of nodes are active during this period (83%, 71%, 75%, and 71% respectively for call, SMS, conference, and email data) and there are enough events for the dynamics to advance. We calculate the global prevalence of adoption at the end of the event list $t = T$ as the ratio of the number of adopted nodes I to the total number of active nodes N_{act} that participated in at least one event during the run. N_{act} practically equals network size, although some rare nodes may only be active before the random initial event. We obtain the mean prevalence by starting the adoption process from 10^3 initial conditions for the larger call and SMS data sets, while for the smaller email and conference data sets we used 10^4 initial conditions.

TABLE II. Network correlations that are preserved (✓) or destroyed (x) by the two reference models.

Correlations	Orig.	RTS	RO
Network structure	✓	✓	✓
Daily patterns	✓	✓	x
Single link event correlations	✓	x	✓
Link-link event correlations	✓	x	x

III. RESULTS

A. Stochastic model

Prevalence and memory length. Figure 3 displays the outcomes of the stochastic model for all four data sets, in terms of the average final fraction of adopters as a function of the memory parameter τ , i.e., the length of the time window within which contacts from adopted individuals are counted. Overall, it is seen that, with the shortest time windows, prevalence is low. When τ is increased, there is a transition regime where prevalence grows, followed by a plateau where increasing τ has only small effects. The low- τ behavior is as one might reasonably expect: increasing memory length means that there are more chances for contacts with adopters. Note that the contrary is observed in the model of Ref. [14], where the number of contacts with adopters is normalized by the total number of contacts within the time window, and because of this, longer memory results in lower adoption rates. It is also of interest to look at the position of the transition region, as it defines the time scale for adoption. In the call network, having a memory τ less than one hour does not facilitate adoption, whereas increasing the memory to more than one week becomes redundant. In the conference network, τ longer than one day becomes redundant.

It is also clear that there are major differences between the data sets: for high enough values of τ , the conference network has the highest prevalence, followed by the call network, whereas the fraction of adopters in the SMS and email networks always remains very low. We have tested that the fraction of adopters in the SMS network is low even when

FIG. 3. (Color online) Average final fraction of adopters in the stochastic model as a function of memory length τ for (a) call, (b) SMS, (c) conference, and (d) email data sets. Curves are shown for the original event sequence (○) and the two null models: random time shuffling (RTS) (□) and random offset (RO) (★). For all networks except SMS, there is a transition region where prevalence increases with memory length and then saturates. The prevalence is much lower for the SMS and email networks as compared to the call and conference networks. The RTS model facilitates adoption dynamics, whereas the RO model hinders adoption. Shades of gray represent memory lengths of one hour, one day, and one week. Standard errors are smaller than symbol size.

the directionality of the events is ignored. This indicates that the reason for the low fraction of adopters lies in the temporal and topological features of the networks rather than in the directionality of the links: the main reasons are the low link density of the network, the existence of a large number of links with relatively low contact frequency, and the fact that conversations by text messages give rise to trains of repeated events only between the same pairs of nodes. These factors lower the likelihood of having temporal paths via which the diffusion can spread [23]. In addition, it is unlikely that nodes whose links have only few events receive contacts from multiple sources, and because of this the probability of crossing the adoption threshold remains low.

Effects of temporal correlations. In order to understand the effects of various kinds of temporal correlations on the adoption dynamics, we next focus on the effects of the reference models on the outcome. The average prevalence as a function of τ for the two reference models, random time shuffle (RTS) and random offset (RO), are also displayed in Fig. 3.

Let us focus on the RTS reference model, and the two networks (conference, calls) where global cascades take place and the effects of reference models are best visible. Contrary to the absolute threshold model in Ref. [14] and the model in Ref. [15], the RTS procedure that removes burstiness from the event sequences yields networks where the adoption rate is increased. This major difference to the purely temporal threshold models has to do with the effects of adoption from multiple neighbors. In the aforementioned models, bursts of contacts from a single neighbor can drive a node above the adoption threshold—these bursts that facilitate spreading are destroyed by the RTS model. However, in our case, contacts from multiple adopted neighbors drive the adoption. As the RTS model spreads out bursty sequences of contacts more evenly in time, the likelihood of contacts with multiple neighbors in a given time window is increased. In addition, the RTS model also increases the probability of having a time-respecting path between any two nodes at any given point of time, hence facilitating adoption dynamics especially in the short memory region where multiple contacts are unlikely even with the reference model.

This picture becomes more complete with the results of the RO model. Compared to the Poissonian event sequences of the RTS model, burstiness of contacts in the original event sequences clusters them in time and decreases the likelihood of contacts with multiple neighbors within short time windows. However, if there are positive timing correlations between the bursty contact patterns of the links of a node, they should facilitate adoption. This is exactly what is seen in the prevalence curves from the RO model (Fig. 3) that randomly shifts the bursty event sequences on links, destroying timing correlations between adjacent links while retaining burstiness on individual links. For both the call and conference networks, applying the random offset results in a decrease of adoption prevalence. Hence one can conclude that, in the original event sequences for the call and conference networks, there are correlations between contact timings of adjacent links that facilitate adoption. At the same time, for the SMS and email networks, global cascades do not take place even for any of the reference models.

B. Deterministic model

Effects of f . Next, we consider the deterministic version of the threshold model, where adoption happens if and only if a fraction f of the node's network neighbors are adopters and in contact with the node within the time period τ . Hence we can control the ease of adoption propagation by changing the threshold fraction f . When f is very low, $f \sim 1/k_{\max}$, where $k_{\max}$ denotes the maximum degree of the network, all events between adopted and susceptible nodes lead to an adoption. In this case, the model's dynamics are similar to the SI model. For any value of f , contact with a single adopter is sufficient for triggering adoption of a node if the node's degree $k \leq 1/f$. On the other hand, for large values of the threshold f , propagation of adoption would require that almost the whole neighborhood of a node are adopters and in contact with the node within the memory period τ . This means adoption typically cannot propagate for large values of f .

Figure 4 displays the fraction of adopters in the deterministic model as a function of f , together with the fraction of nodes with $k \leq 1/f$ that obeys SI dynamics, N_{SI}/N . In these simulations, the memory is fixed to $\tau = \infty$, so that the effect of f can be determined in isolation. In all cases, the values of f need to be small in order for the adoption to take off. For instance, in the call network adoption does not propagate at all if $f \geq 0.12$. This is noteworthy, since at this threshold fraction more than 80% of the nodes follow SI dynamics and contact with a single adopter is sufficient for adoption. Only nodes with $k > 8$ obey the true threshold mechanism. Because the majority of nodes follow SI dynamics and because pure SI dynamics would always infect the whole

FIG. 4. (Color online) Average final fraction of adopters in the deterministic model as a function of fraction f with $\tau = \infty$ for the (a) call, (b) SMS, (c) conference, and (d) email data sets. The adoption process is shown for the original event sequence ($\circ$) and the two null models: (i) random time shuffle (RTS) ($\square$) and (ii) random offset (RO) ($\star$). The black solid line represents the fraction of nodes in a network that obey pure SI dynamics at given f value. Except for the conference data, the majority of the nodes must obey SI dynamics in order for a global diffusion to happen. Standard errors are smaller than symbol size.

FIG. 5. (Color online) Average final fraction of adopters in the deterministic model as a function of memory τ for the (a) call, (b) SMS, (c) conference, and (d) email data sets. Again, the adoption process is shown for the original event sequence ($\circ$) and the two null models: (i) random time shuffle (RTS) ($\square$) and (ii) random offset (RO) ($\star$). The fraction f is now chosen so that the curves have a transition as τ changes. The difference in the fraction of adopters in the RTS and RO null model is more evident in this case. As the temporal path does not exist between all the nodes, the maximum prevalence is also shown for all the data sets for RTS null model. Standard errors are smaller than symbol size.

network, this indicates that high-degree nodes typically block the propagation of adoption, as they are unlikely to interact with enough adopters. For the SMS network, the adoption propagates only when $f \leq 0.08$, with more than 95% of the nodes following SI dynamics. In contrast, in the email network global adoption occurs at $f \leq 0.04$, with about 80% of nodes following SI dynamics. The differences imply that interactions with multiple neighbors are more common in the email network as compared to the SMS network. Finally, for the conference network, global adoption arises even when most of the nodes follow pure threshold dynamics. This is because of the high frequency of interaction between nodes.

Effects of memory length τ . Next, we study the behavior of the system as a function of the memory τ . The results are shown in Fig. 5, where the fraction f for each data set is chosen so that the adoption can propagate, given high enough τ . Unlike with the stochastic model, even for the SMS and email networks, a substantial fraction of nodes become adopters in the end. As discussed above, this occurs at low thresholds, and most nodes obey SI dynamics, i.e., a single contact with an adopted neighbor triggers adoption. This is also reflected in the observation that adoption prevalence in the SMS and email networks is barely affected by the memory τ . For the call and conference networks the outcomes are rather similar to the stochastic model. The only difference is that the transition region where prevalence increases is associated with lower values of τ . In the call network, this happens with $\tau \sim 1$ d, and for the conference network, $\tau \sim 1$ h. For the stochastic model, the corresponding τ values are $\tau \sim 1$ week for calls and $\tau \sim 1$ d for the conference.

Effects of temporal correlations. The effects of the reference models for the deterministic model are consistent with the outcomes in the stochastic case. As before, for all f the fraction of adopters increases for the RTS model and decreases for the RO model in all data sets. However, the deterministic case reveals larger differences between the reference models, especially in the low τ region. This is because of the altered temporal path structure and its effects on the SI dynamics that most nodes obey.

C. Measuring temporal correlations

Density of preceding events. The above results point out that there are correlations between the timings of events on adjacent links, and their effects are especially strong for the call network. The existence of such correlations is also known from earlier studies on, e.g., temporal motifs [21,22] and interevent time distributions [11]. In order to directly measure these correlations and to look for possible characteristic time scales, we compute the density of preceding events (or event-triggered correlation function [24,25]). The target is to study what happens just before nodes participate in events. For directed events, we define the density of preceding events as follows: for every outgoing event of a node, the time differences Δt to all its earlier incoming events are calculated (see Fig. 6). The probability density functions for the Δt values for all nodes and events then represent the average rate of incoming events preceding an outgoing event, and any timing correlations should be visible in this PDF. Especially, if there is a characteristic time for the correlations of events on adjacent links, this should be visible as a peak in the PDF. For undirected events, this PDF reduces to the usual distribution of interevent times. For both undirected and directed events, we repeat the procedure separately for two cases: (i) where preceding events with all neighboring nodes are taken into account (case *all*), and (ii) for three-node chains of events only, where Δt 's are computed only for those preceding events that do not involve the other party of the trigger event (case *forward*). Thus the difference between the *all* and *forward* cases is that, in the latter, the effect of event cycles of length two, i.e., returned events, is removed.

FIG. 6. (Color online) Calculation of the density of preceding events, for directed events. The red arrows represent the events; event times are indicated next to the arrows. All outgoing events from a node act as triggers for calculating time differences to preceding incoming events. The set of time differences for all nodes and all trigger events is then used to construct a probability density function.

FIG. 7. (Color online) Density of preceding events (top row) and fraction of accessible neighborhood in a given time window (bottom row) for the call and SMS data. Triangles and crosses indicate preceding event densities for all events and forward chains (events excluding the neighbor participating in the trigger event), respectively. (a) For both cases, for the call network there is a peak in the distribution at around 25 s. (b) For the SMS data, the peak is only evident when all earlier events are considered—as the SMS events are directed, this means that the peak in the preceding event density is mainly due to returned text messages. (c), (d) The complementary cumulative distribution of $\langle k_i^\tau \rangle / k_i$ for the call and SMS networks, respectively. $\langle k_i^\tau \rangle / k_i = 0$ is not shown because the axis is logarithmic; the fraction of nodes with zero contacts is reflected in the starting point of the curves. For any given time window the accessible neighborhood for the SMS network is much smaller than for the call network. The dashed lines indicate 50% for both axes. The hourly CDF's are first calculated independently for each hour of the day and then averaged; the shaded areas represent the 75% confidence interval for the mean of the 24 different CDF's. For longer memory intervals, the 75% confidence interval for the mean is smaller than the linewidth.

The preceding event densities for the call and SMS data sets are shown in Fig. 7. There is a characteristic peak for event correlations on adjacent links: the PDF peaks at around ~ 25 s for the call data, and at around ~ 45 s for the SMS data when all events are considered. For forward event chains (events excluding the neighbor of the trigger event), the call network peak shifts to ~ 80 s. In both cases, the time scales are fairly short, of the order of one minute, and thus event-event correlations have effects on the threshold model for a wide range of memory parameter values. For the directed text message network, there is no clear peak if only forward chains are counted. This is likely because text messages are typically part of a conversation, where two individuals repeatedly exchange multiple messages. These are redundant for adoption dynamics.

Fraction of contacted neighbors and memory length. In order to better understand the effects of the memory length on the two threshold models (stochastic, deterministic), it makes sense to directly measure how the amount of interactions

with the neighborhood of a node depends on the memory window size. We count the number of nodes the focal node interacts with in a given window of size τ as $k_i[t, t + \tau] \equiv k_i^\tau$ and divide this number by the degree of the node in the aggregated static network, $k_i[0, T] \equiv k_i$. Averaging over nodes and windows yields the typical fraction of neighbors a node connects with within a given memory length τ . Note that, for the stochastic threshold model, the fraction $\langle k_i^\tau \rangle / k_i$ corresponds to the probability of adoption $\phi(i, \tau)$ when the whole neighborhood of a node has already adopted. In Fig. 7 we show the complementary cumulative distribution of $\langle k_i^\tau \rangle / k_i$ for the call and SMS networks for different memory lengths $\tau = (\text{hour, day, week, the whole data period } T)$. For the call network, it is clear why the memory window size has to be of the order of days for the adoption to propagate (Figs. 3 and 5): for memory of one hour, the fraction of contacted neighbors is very low. For the SMS network, even for a window of one week, $\sim 60\%$ of the nodes have $\langle k_i^\tau \rangle / k_i$ less than 0.05, indicating rather infrequent communication.

IV. DISCUSSION

In this paper, we have introduced stochastic and deterministic versions of the topological-temporal threshold model of adoption, and studied their behavior using empirical temporal network data sets as substrates for the threshold dynamics. These models can be argued to reflect social diffusion processes better than some models proposed earlier [14,15], since adoption probability directly depends on the number of contacted nodes, a mechanism observed in real-world experiments [6]. However, in addition to this, an equally important factor of motivation is that these models can be used as probes of temporal network structure: because of their design, our models are sensitive to timing correlations of the event trains on links.

For the stochastic threshold model, the call and conference networks allow global cascades for large enough memory windows (days for calls, hours for the conference network), whereas the SMS and email networks did not support such cascades because of the sparsity of interactions and lack of timing correlations between contacts on adjacent links. The characteristic window sizes are meaningful for the networks in question: in the face-to-face network recorded in a conference setting, people participate and switch between conversations over short time scales, whereas interactions via calls are less frequent. For the deterministic model, where adoption only takes place when the fraction of adopters in a node's neighborhood exceeds a fixed threshold fraction, global adoption only took place when the threshold fractions were set to low enough values and most nodes effectively followed SI dynamics. This is consistent with the results of threshold dynamics on static networks, where low-degree vulnerable nodes drive the adoption dynamics [9]. One reason for this comes directly from model design: the higher the degree of a node, the larger the number of neighbors that first need to adopt and subsequently be in contact with the node within the memory window. Thus the effects of hubs are very different from ordinary (SI, SIR) spreading processes: while they may act as superspreaders once they have adopted, making them adopt is very difficult.

Randomly shuffling the times of all events in the empirical networks was seen to facilitate the adoption process and results in larger prevalence, similarly to earlier observations with the SI model [11]. Thus, generally speaking, the burstiness of contacts that is ubiquitous in temporal networks hinders adoption because of increased waiting times on links and redundant repeated events. The time shuffling procedure destroys this burstiness, spreading events more evenly across time and giving rise to an increased number of temporal paths ending at nodes within short time windows. Note that this result differs from threshold models that do not explicitly require contacts from multiple adopted neighbors. However, there is a clear competing effect arising from correlations between contact trains. Randomly time shifting the contact events on links was seen to decrease prevalence of adoption. This procedure destroys all correlations between timings of contacts on adjacent links. Because such correlations by design facilitate adoption in our models, the clear decrease in prevalence after randomly offsetting contact trains indicates that such correlations are abundant in the studied networks, as also revealed by direct measurements.

Some features of the dynamics of our threshold models are similar to the behavior of SI and SIR models on temporal networks [11,13]. For example, burstiness hinders the speed

and fraction of adopters in both cases, because it increases waiting times along temporal paths. However, there are also noticeable differences. The effect of timing correlations is much more evident for the threshold models, as the threshold dynamics is driven by multiple contacts within a short time window. Further, the sparsity of events and lack of timing correlations may block the adoption process, as seen for the stochastic model applied to the SMS and email networks. For the deterministic model, the adoption in these networks occurs only for very low threshold values, and even then the fraction of adopters is independent of the memory length. Both features are very different from the SI and SIR models, where either adoption always occurs or the final fraction of adopters changes with the model parameters (for SIR, the basic reproduction number). Overall, threshold models can be seen as an addition to the family of models of contagion in temporal networks.

ACKNOWLEDGMENTS

We acknowledge financial support by the Academy of Finland, Project No. 260427. We thank A.-L. Barabási for the call and SMS data.

-
- [1] R. B. Cialdini and N. J. Goldstein, *Annu. Rev. Psychol.* **55**, 591 (2004).
 - [2] M. J. Salganik, P. S. Dodds, and D. J. Watts, *Science* **311**, 854 (2006).
 - [3] E. Bakshy, I. Rosenn, C. Marlow, and L. Adamic, in *Proceedings of the 21st International Conference on World Wide Web* (ACM, New York, 2012), pp. 519–528.
 - [4] C. Castellano, S. Fortunato, and V. Loreto, *Rev. Mod. Phys.* **81**, 591 (2009).
 - [5] R. M. Anderson and R. M. May, *Infectious Diseases of Humans: Dynamics and Control*, 1st ed. (Oxford University Press, New York, 1992).
 - [6] D. Centola, *Science* **329**, 1194 (2010).
 - [7] J. Ugander, L. Backstrom, C. Marlow, and J. Kleinberg, *Proc. Natl. Acad. Sci. U.S.A.* **109**, 5962 (2012).
 - [8] M. Granovetter, *Am. J. Sociol.* **83**, 1420 (1978).
 - [9] D. J. Watts, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 5766 (2002).
 - [10] P. Holme and J. Saramäki, *Phys. Rep.* **519**, 97 (2012).
 - [11] M. Karsai, M. Kivelä, R. K. Pan, K. Kaski, J. Kertész, A.-L. Barabási, and J. Saramäki, *Phys. Rev. E* **83**, 025102 (2011).
 - [12] M. Kivelä, R. K. Pan, K. Kaski, J. Kertész, J. Saramäki, and M. Karsai, *J. Stat. Mech.* (2012) P03005.
 - [13] G. Miritello, E. Moro, and R. Lara, *Phys. Rev. E* **83**, 045102 (2011).
 - [14] F. Karimi and P. Holme, *Physica A* **392**, 3476 (2013).
 - [15] T. Takaguchi, N. Masuda, and P. Holme, *PLoS ONE* **8**, e68629 (2013).
 - [16] P. S. Dodds and D. J. Watts, *Phys. Rev. Lett.* **92**, 218701 (2004).
 - [17] J.-P. Eckmann, E. Moses, and D. Sergi, *Proc. Natl. Acad. Sci. U.S.A.* **101**, 14333 (2004).
 - [18] L. Isella, J. Stehlé, A. Barrat, C. Cattuto, J.-F. Pinton, and W. Van den Broeck, *J. Theor. Biol.* **271**, 166 (2011).
 - [19] SocioPatterns, <http://www.sociopatterns.org/>.
 - [20] P. Holme, *Phys. Rev. E* **71**, 046119 (2005).
 - [21] L. Kovanen, M. Karsai, K. Kaski, J. Kertész, and J. Saramäki, *J. Stat. Mech.* (2011) P11005.
 - [22] L. Kovanen, K. Kaski, J. Kertész, and J. Saramäki, *Proc. Natl. Acad. Sci. U.S.A.* **110**, 18070 (2013).
 - [23] R. K. Pan and J. Saramäki, *Phys. Rev. E* **84**, 016105 (2011).
 - [24] L. Kovanen, Master's thesis, Aalto University, Finland, 2009.
 - [25] F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek, *Spikes: Exploring the Neural Code* (MIT Press, Cambridge, MA, 1999).