

Supporting Information

Kovanen et al. 10.1073/pnas.1307941110

SI Text

Analyzing Data with Null Models

After counting the occurrence of motifs in empirical data, a similar problem is encountered with both static network motifs (1) and temporal motifs (2): To say whether a motif count is high or low, we need to compare it to something. Because typically no obvious reference is available, it has become common to construct one artificially by randomizing some aspects of the empirical data. This kind of reference is commonly known as null models.

Network Motifs and the Configuration Model. Before going into temporal motifs, it is useful to understand how null models have been used to analyze static network motifs. The problems that arise in this context are very similar to those we face when developing null models for temporal motifs.

The most widely used null model for analyzing network motifs is no doubt the configuration model, a random network conditional on node degrees (3). [To be precise, Milo et al. (1) originally used a slightly different null model: They created a random network conditional on the occurrence of smaller motifs than those being studied. For the purposes of our discussion here, this detail is, however, not relevant.] Milo et al. (1) suggested that the motifs that are “likely to be important” can be identified by comparing the empirical count $C(m)$ to that in the null model, $\tilde{C}(m)$. In fact, the term “network motif” was originally used to refer to only those patterns whose number is so high that it is unlikely to occur in the null model. In 2004, Artzy-Randrup et al. (4) pointed out several problems in the use of configuration model—or in fact any random network—to define important motifs. They constructed a simple random network on a 30×30 lattice where nearby nodes are connected with higher probability, and showed that the same motifs are overrepresented in this network as in the neural network of *Caenorhabditis elegans* that was originally studied by Milo et al. Comparison with the configuration model cannot reveal whether the occurrence of motifs is due to randomness, but only whether the empirical network was generated by the configuration model.

The logic is made more transparent by writing out the null hypothesis being tested. Because the only thing the configuration model retains is the node degrees, comparing motif counts against those in the configuration model corresponds to testing the null hypothesis:

H_0 : Motif count depends only on node degrees.

If we find motif counts that are unlikely to occur in the data, we can reject the null hypothesis: Motif counts are not explained by the node degrees alone.

Of course, the result is nearly always trivial: It is hardly surprising to find that an empirical network has more structure than the configuration model. Typically, this is also not the result we set out to obtain. What we want instead is to identify those motifs whose occurrence is not explained by the configuration model. However, comparison against the configuration model does not allow such inference. If the count of even a single motif differs sufficiently from that expected in the configuration model, H_0 is very likely to be false. In addition, if H_0 is false, it cannot explain the occurrence of any motif, not even of those whose occurrence happens to agree with that observed in the configuration model. (There are often some motifs whose count is similar to that in the configuration model. This is typically due to some kind of conservation law: If some motifs are overrepresented, others are

necessarily underrepresented, and some will happen to be approximately equally common.)

Even if we cannot explain the occurrence of motifs, we might still be able to obtain useful information by studying the magnitude of deviation from the null model. Any difference observed between $C(m)$ and $\tilde{C}(m)$ is obviously due to differences between the empirical data and the null model, and thus to interpret this deviation we need to know what this difference is. In the case of the configuration model, for example, this difference consists of the network structure destroyed by randomizing edges. However, finding out that some motif is, say, 50% more common because of the network structure other than node degrees is a very broad statement. To know more precisely what is causing the difference, we need an explicit list of features destroyed by the configuration model: correlations between degrees of neighboring nodes, number of triangles around the nodes, distribution of shortest path lengths, etc. However, as such list would allow reconstructing the empirical network, it would constitute a perfect network model and is obviously nontrivial to construct. In addition, even if we did succeed, what is the value of knowing that a motif is more common because of all of these features? Is that really what we set out to find? Furthermore, we only observe the total effect of all removed features. Some of these features might even have an opposite effect to motif counts.

The value of using null models is not in disproving the null hypothesis or in finding motifs that are “explained” by the null model, but in being able to gain information from the deviation between $C(m)$ and $\tilde{C}(m)$. Because the meaning of this deviation comes from the difference between the empirical data and the null model, the null model should be constructed so that this difference is explicitly known and matches the research question.

Null Models for Temporal Motifs. To illustrate the use of null models for analyzing temporal motifs, we study a series of synthetic datasets. All datasets have $N = 10,000$ nodes assigned into two classes, red (r) and blue (b). In all datasets, we generate events until there are on average 100 events per edge. We will see that, as more features are added into the synthetic data, the null models fail when the data no longer satisfy the corresponding null hypothesis. We compare the following three null models:

Color-shuffled reference. Color-shuffled reference is constructed by shuffling the colors of all nodes. Because this null model retains the original event data except for the assignment of node colors, the null hypothesis being tested is as follows:

H_0^C : Motif counts depend on node colors only via the number of nodes of each color.

As with the configuration model above, it is not easy to list explicitly the features that are destroyed by this null model. Such list would contain all possible ways two nodes can differ in a temporal network, including preferred connectivity, weights of adjacent edges, and temporal behavior in general.

Time-shuffled reference. Time-shuffled reference is constructed by shuffling the time stamps of events. Because this null model retains the aggregate network and the number of events on each edge, the null hypothesis being tested is as follows:

H_0^T : Motif counts do not depend on temporal correlations between events.

Because this null hypothesis states explicitly what is being removed, it provides an interpretation for the deviation from the null model: Comparison against this reference reveals that

a motif is more or less common because of temporal correlations. Whether this information is useful is another matter. “Temporal correlations” is a very broad notion, comprising the temporal distribution of events on a single edge as well as the correlations of events on neighboring edges.

Color-equivalent reference. Color-equivalent reference is the one described and used in the main text. Because this null model removes differences in temporal behavior between node colors, the null hypothesis is as follows:

H_0^C : Motif counts do not depend on differences in temporal behavior between node colors.

As with the time-shuffled reference above, this null hypothesis states explicitly what is being removed, allowing us to give a meaning for the difference between $C(m)$ and $\tilde{C}(m)$: This difference reveals how much more or less common a motif is because of differences in the temporal behavior of node colors, or in other words, those aspects of node colors not observable in the weighted aggregate network.

We begin with a simple synthetic data where the two node colors have different cardinality: A node is colored red with a probability of 40%; otherwise, the node is blue. To generate the event data, we first create an underlying network by connecting nodes at random with probability $P=0.001$. Events are then generated so that at any time step an event starts on a given edge with probability $q=0.0001$, assuming neither node is currently involved in an event. The duration of the events is drawn from a geometric distribution with mean 10.

For these simple data, all three null hypotheses are true, and, as expected, the z -score distributions shown in Fig. S1 are all concentrated around zero. Only a small change is needed to make H_0^C false; this happens for example if the edges of the underlying network are not independent of node colors. To illustrate, Fig. S2A shows the z -score distributions when the probability of an edge between nodes depends on node colors so that $P_{rr}=0.0012$, $P_{rb}=0.0008$, and $P_{bb}=0.001$. A similar result is obtained if the event probabilities depend on node colors. Fig. S2B shows the result when the underlying network is uniform but the event probabilities between different colors are $q_{rr}=0.00012$, $q_{rb}=0.00008$, and $q_{bb}=0.0001$. The other two null hypotheses, however, continue to hold: In both cases, the events are uncorrelated and there are no temporal differences between node colors.

In empirical datasets, the events are practically never independent of each other. To illustrate the effect this has on the null model analysis, we add correlations to our synthetic data. Random events occur as before, but in addition triggered events may take place on neighboring edges after each random event. The probability of these triggered events decreases exponentially with time, with parameters selected so that with probability $P_\tau=0.02$ at least one triggered event occurs on a neighboring edge during the following $\tau=50$ time steps. As shown by Fig. S2C, H_0^T is now false but both H_0^C and H_0^E are true.

Combining all of three features—homophily, different event probabilities, and correlated events—produces data where both H_0^C and H_0^T are false even though there are still no temporal differences between node colors (Fig. S3A). To see how the null models perform when there are temporal differences between node colors, we increase the occurrence of some causal two chains exactly as done in the main text: If a random event occurs between two nodes of the same color, for the following $T_t=100$ time steps the recipient has an additional probability of $P_t=0.001$ to initiate an additional triggered event toward a random neighbor other than the source of the original random event.

The resulting z -score distributions are shown in Fig. S3B. The color-shuffled reference is clearly unable to detect the triggered two chains. The time-shuffled reference is slightly better, as the z scores for these motifs are even higher than for other motifs. Unfortunately, real empirical data rarely have such pronounced

temporal differences; typically, they would be lost among all other motifs that are still significantly overrepresented.

The z -score distribution for the color-equivalent reference has three peaks. Although most motifs show no difference ($z \approx 0$), a small number of motifs have $z \gg 0$ and a similar number has $z \ll 0$. To see what these motifs are, Fig. S4 shows all two-event temporal motifs ordered by the ratio $r(m)$ corresponding to this null model. As expected, the motifs with $z \gg 0$ are causal two chains where the first event takes place between nodes of the same color. In addition, because we measure the relative occurrence, the motifs with $z \ll 0$ are those causal two chains where the first event takes place between nodes of different colors.

A closer look at Fig. S4 reveals that only three of the causal two chains have $z \gg 0$. The fourth causal two chain expected to be overrepresented, the one with three red nodes, has $r(m)=0.98$ and is in fact less common than in the reference. This reason is that the object of comparison, $\tilde{C}(m)$, is based on average uncolored motif counts. Because the three most common causal two chains have very high occurrence—due to blue nodes being more common—they hike up the value of $\tilde{C}(m)$ to the extent that it becomes larger than the actual motif count of the all-red causal two chain. [The all-red causal two chain is still overrepresented if we compare the absolute count $C(m)$ in these data to that in data that are otherwise identical but lack the preferential triggering.] Another way to understand this is to think of $\tilde{C}(m)$ as the weighted average of colored motif counts, with weighting done by the entire structure of the aggregate network. Increasing the motif count of some colored variants will necessarily reduce the relative occurrence of others.

Proof That $\tilde{C}(m)$ Is an Unbiased Estimate

To see that $\tilde{C}(m)$ is an unbiased estimate of $C(m)$ when H_0 is true, we write its expected value as follows:

$$\begin{aligned} E[\tilde{C}(m)] &= \sum_{\ell|m_\ell=m} E[\tilde{C}_\ell(m)] \\ &= \sum_{\mathbf{w}} \sum_{\ell|m_\ell=m, \mathbf{w}_\ell=\mathbf{w}} E[\tilde{C}_\ell(m)] \\ &= \sum_{\mathbf{w}} |\{\ell|m_\ell=m, \mathbf{w}_\ell=\mathbf{w}\}| \cdot E[\tilde{C}_\ell(m)|\mathbf{w}_\ell=\mathbf{w}]. \end{aligned}$$

However, the empirical motif count can be written as follows:

$$\begin{aligned} C(m) &= \sum_{\ell|m_\ell=m} C_\ell(m) \\ &= \sum_{\mathbf{w}} \sum_{\ell|m_\ell=m, \mathbf{w}_\ell=\mathbf{w}} C_\ell(m) \\ &= \sum_{\mathbf{w}} |\{\ell|m_\ell=m, \mathbf{w}_\ell=\mathbf{w}\}| \cdot \bar{C}_{\mathbf{w}}(m), \end{aligned}$$

where $\bar{C}_{\mathbf{w}}(m)$ is the average count of motif m at locations with weight sequence \mathbf{w} . Now $E[\tilde{C}(m)]=C(m)$ if $E[\tilde{C}_\ell(m)|\mathbf{w}_\ell=\mathbf{w}]=\bar{C}_{\mathbf{w}}(m) \forall \mathbf{w}$, which is indeed the case when H_0 is true.

Generating Samples from the Null Model

Fig. S5 presents an algorithm that can be used to generate samples from the null model described in *Materials and Methods*. The input consists of the temporal network (event set E), the maximum number of events in temporal motifs (n_{\max}), number of independent samples to generate (N_{samples}), and the time window Δt used to identify temporal motifs. The algorithm outputs motif counts in the empirical data (C) and the motif counts sampled from null model (\tilde{C}), corresponding to the count of motif m under the null hypothesis that motif counts do not depend on node types, given the structure of the weighted aggregate network.

The algorithm consists of three parts. On lines 2–8, we identify all temporal motifs in the empirical data and count the occurrence of motifs by location. This results in two data structures: $M[\ell]$ gives the number of temporal motifs at location ℓ (note that the corresponding motif m_ℓ is uniquely defined by ℓ), and $C[m]$ gives the total count of motif m in the empirical data.

On lines 9–18, these two data structures are used to construct the distributions $P(m^*, \mathbf{w})$ for all combinations of the untyped motif m^* and topology \mathbf{w} . The result is a data structure P such that $P[(m^*, \mathbf{w})][C]$ is the number of locations with weight sequence \mathbf{w} and topology defined by m^* that have exactly C occurrences of the motif m^* . Note that node colors are not used during this step; the distributions are constructed under the null hypothesis that node colors have no effect, in which case the distributions would be identical for different node colors.

On lines 20–26, we draw samples from $P(m^*, \mathbf{w})$. We again go through all locations, but this time take into account the node colors. Because the final sample $\tilde{C}(m)$ is a sum over only those locations that have correct node colors, the structure of the aggregate network is taken into account. Note that it is relatively cheap to generate multiple samples for $\tilde{C}(m)$ (lines 24–26); the

only thing that needs to be repeated is drawing samples from $P(m^*, \mathbf{w})$.

Homophily and Sex Differences for Calls and SMS

Tables S1 and S2 show full results on temporal homophily for both calls and SMS. Parts of Table S1 and S2 are presented in the main text.

Sensitivity of Results to the Choice of Δt

The exact value of the time window used in the analysis, $\Delta t = 10$ min, is of course arbitrary. As discussed in ref. 5 in the context of aggregation of temporal networks, selecting the proper length for a time window is a balancing act. Low values of Δt decrease motif counts and therefore increase noise, whereas high values of Δt smooth the results as unrelated events are more likely to be connected.

In our case, however, the ratios $r(m)$ are not sensitive to the exact value of Δt , as shown in Fig. S6: The difference the values of $r(m)$ between $\Delta t = 10$ min and $\Delta t = 1$ h is still comparable to the difference between 2 consecutive months calculated with $\Delta t = 10$ min.

1. Milo R, et al. (2002) Network motifs: Simple building blocks of complex networks. *Science* 298(5594):824–827.
2. Kovanen L, Karsai M, Kaski K, Kertész J, Saramäki J (2011) Temporal motifs in time-dependent networks. *J Stat Mech* 2011(11):P11005.
3. Newman MEJ, Strogatz SH, Watts DJ (2001) Random graphs with arbitrary degree distributions and their applications. *Phys Rev E Stat Nonlin Soft Matter Phys* 64(2 Pt 2):026118.
4. Artzy-Randrup Y, Fleishman SJ, Ben-Tal N, Stone L (2004) Comment on “Network motifs: Simple building blocks of complex networks” and “Superfamilies of evolved and designed networks”. *Science* 305(5687):1107, author reply 1107.
5. Sulo R, Berger-Wolf T, Grossman R (2010) Meaningful selection of temporal resolution for dynamic networks. *Proceedings of the Eighth Workshop on Mining and Learning with Graphs* (ACM, New York), pp 127–136.




Fig. S1. The distribution of z scores of all colored two-event motifs, presented separately for the three null models. The gray line shows the Gaussian distribution with zero mean and unit variance for reference. The data have been artificially generated as follows. We first create an undirected, unweighted underlying network with $N = 10,000$ nodes. The nodes are randomly assigned one of two colors: red with probability of 40% and blue otherwise. An edge is created between any two nodes with probability $P = 0.001$. Events are uncorrelated: An event occurs on an edge with probability $q = 0.0001$ per times step (assuming neither node is currently involved in an event) with event duration drawn from a geometric distribution with mean 10. Events are generated from this process until there are on average 100 events per edge. Two-event temporal motifs are then identified with time window $\Delta t = 100$. The distributions shown are averages of 20 independent datasets; for each dataset, we generate 20 values from each null model to calculate $\mu(\tilde{C}(m))$ and $\sigma(\tilde{C}(m))$. Because for these simple data all three null hypotheses are true, the z-score distributions are centered around zero.

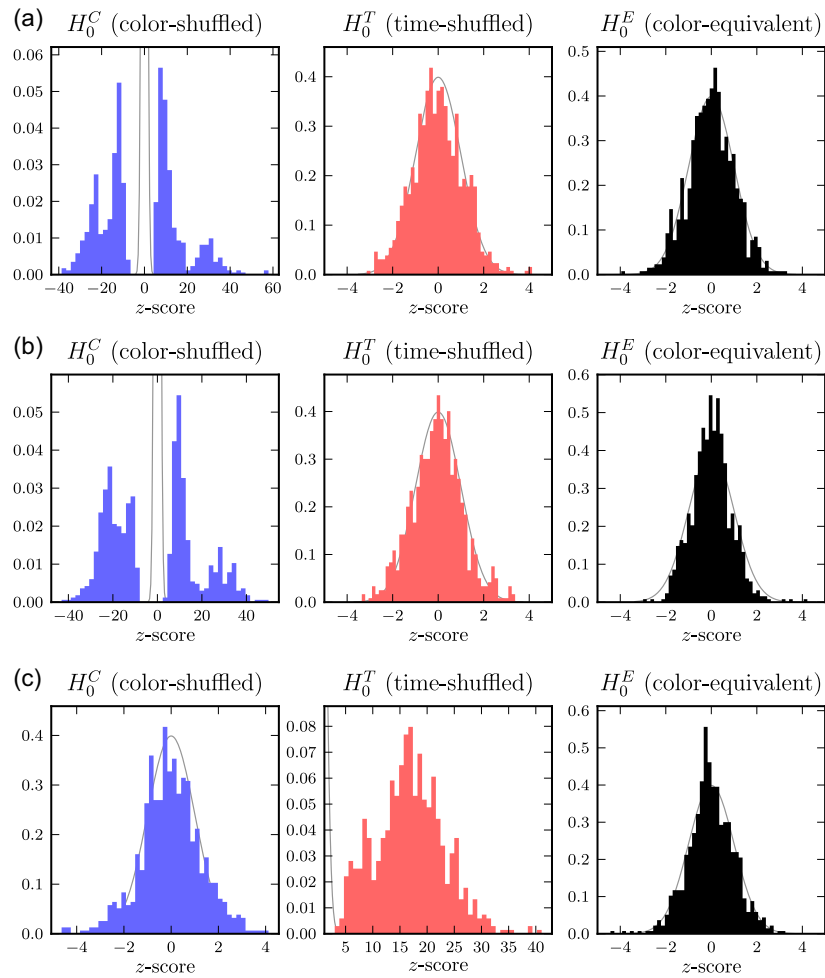


Fig. S2. Three variations of the synthetic data used in Fig. S1. (A) The underlying network is modified so that edges are preferentially created between nodes of the same color. H_0^C is false because the number of nodes of each color is no longer sufficient to explain the occurrence of motifs. (B) The event probabilities are changed so that events are more likely to occur between nodes of the same color. H_0^C is false for the same reason as above. (C) Now the underlying network is uniform and the occurrence of events is independent of node colors, but the events are correlated between neighboring edges. This time H_0^C is false ($z \gg 0$) but H_0^E is true. In all three datasets, H_0^T is true because there are no temporal differences between node colors.

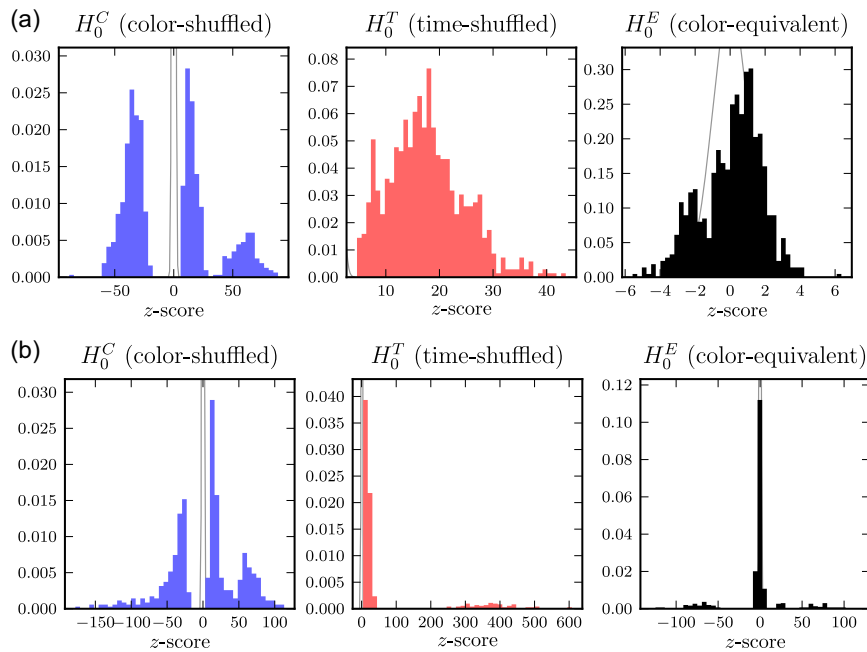


Fig. 53. (A) Combining the three modifications used in Fig. S2 A–C results in both H_0^C and H_0^T being false; H_0^E is still (approximately) true. (B) Results after adding temporal differences to the data used in A. Only H_0^E can distinguish motifs that really exhibit temporal differences ($|z| \gg 0$) from all others ($z \approx 0$).

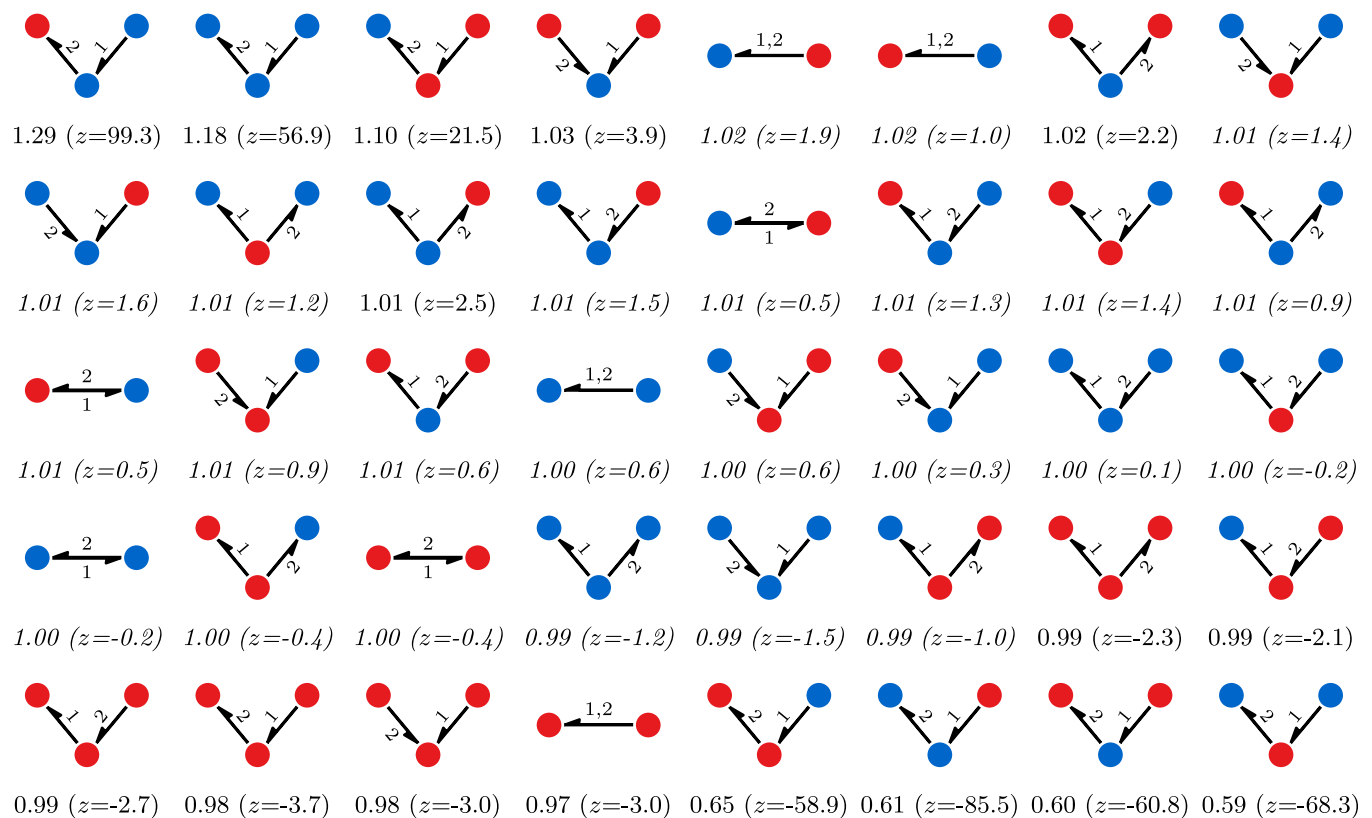


Fig. 54. All two-event motifs corresponding to Fig. S3B, ordered by ratio $r(m) = C(m)/\bar{C}(m)$. Both ratio and z score are shown below each motif. The three most common motifs are causal two chains where the first event takes place between nodes of same color; these have been artificially made more common. However, the fourth motif in this group—causal two chain where all nodes are red—has $r(m) = 0.98$, and is thus less common than in the reference. This low value of $r(m)$ occurs not because $C(m)$ is low—this absolute count is significantly higher than in a similar data without the triggered events—but because $\bar{C}(m)$ is so high. The value of $\bar{C}(m)$ is based on all causal two chains, and this value becomes large because the first three motifs are so common.


```

1: function SAMPLENULLMODEL( $E, n_{\max}, N_{\text{samples}}, \Delta t$ )
2:   Initialize  $M$  as  $\text{Map}(\ell \rightarrow \mathbb{N})$ .
3:   Initialize  $C$  as  $\text{Map}(m \rightarrow \mathbb{N})$ .
4:   for  $G_t$  in  $\text{VALIDSUBGRAPHS}(E, n_{\max}, \Delta t)$  do
5:     Let  $\ell$  be the location of  $G_t$ .
6:     Let  $m$  be the colored motif corresponding to  $G_t$ .
7:     Increment  $M[\ell]$ .
8:     Increment  $C[m]$ .
9:   Construct weighted aggregate network  $G$  from  $E$ .
10:  Initialize  $P$  as  $\text{Map}((m^*, \mathbf{w}) \rightarrow \text{Map}(\mathbb{N} \rightarrow \mathbb{N}))$ 
11:  for  $\ell$  in  $\text{LOCATIONS}(G, n_{\max})$  do
12:    Let  $m_\ell^*$  be the uncolored motif defined by  $\ell$ .
13:    Let  $\mathbf{w}_\ell$  be the weight sequence at  $\ell$  in  $G$ .
14:    if  $\ell \in M$  then
15:      Let  $C_\ell = M[\ell]$ 
16:    else
17:      Let  $C_\ell = 0$ 
18:      Increment  $P[(m_\ell^*, \mathbf{w}_\ell)][C_\ell]$ .
19:  Initialize  $\tilde{C}$  as  $\text{Map}(m \rightarrow \text{List})$ 
20:  for  $\ell$  in  $\text{LOCATIONS}(G, n_{\max})$  do
21:    Let  $m_\ell$  be the colored motif defined by  $\ell$ .
22:    Let  $m_\ell^*$  be the uncolored motif defined by  $\ell$ .
23:    Let  $\mathbf{w}_\ell$  be the weight sequence at  $\ell$  in  $G$ .
24:    for  $i$  in  $1 \dots N_{\text{samples}}$  do
25:      Let  $\tilde{C}_\ell$  be a sample from  $P[(m_\ell^*, \mathbf{w}_\ell)]$ 
26:      Add  $\tilde{C}_\ell$  to  $\tilde{C}[m_\ell][i]$ .
27:  return  $C, \tilde{C}$ 

```

Fig. S5. Algorithm for generating samples from the null model. E is the event set that defines the temporal network, n_{\max} is the maximum number of events in temporal motifs to study, and N_{samples} is the number of independent samples to generate for each $\tilde{C}(m)$. The function VALIDSUBGRAPHS goes through all valid temporal subgraphs with at most n_{\max} events. The function LOCATIONS goes through all locations—ordered sequences of edges such that the underlying graph formed by these edges is connected—in the aggregate network G that have at most n_{\max} edges. The function returns both the empirical motif counts C and the sampled counts \tilde{C} .




Fig. S6. The results are not sensitive to the exact value of Δt . The plot shows the correlation coefficient ρ between $r(m)$ calculated for month 2 data with $\Delta t = 10$ (denoted by the vertical line), and month 1 data with Δt values ranging from 5 s to 4 h. From $\Delta t = 5$ min to $\Delta t = 1$ h the correlation coefficients is close to that of the 2 consecutive months with the same time window. The calculation includes all two-event motifs that occur at least 10 times in every dataset (only 1,214 different motifs satisfy this condition when $\Delta t = 5$ s, compared with 31,816 motifs with $\Delta t = 10$ min).

Table S1. Temporal homophily of calls (upper) and SMS (lower) for different motifs

Communication type	Motif	A	G	P	A \wedge G	A \wedge P	G \wedge P	A \wedge G \wedge P
Call	Repeated contact	1.08, 1.11	1.12, 1.09	1.09, 1.13	1.12, 1.11	1.05, 1.11	1.11, 1.11	1.09, 1.11
	Returned contact	1.04, 1.01	1.02, 1.01	0.98, 1.06	1.06, 1.01	1.00, 1.02	0.98, 1.03	1.02, 1.02
	Noncausal chain	1.05, 1.03	1.05, 1.03	1.05, 1.01	1.11, 1.03	1.08, 1.03	1.06, 1.03	1.18, 1.03
	Causal chain	1.03, 1.02	1.04, 1.02	1.05, 0.98	1.08, 1.02	1.07, 1.02	1.07, 1.01	1.16, 1.02
	Out-star	1.12, 1.03	1.06, 1.03	1.07, 1.01	1.22, 1.04	1.18, 1.04	1.09, 1.03	1.32, 1.04
	In-star	<i>1.09, 1.04</i>	<i>1.07, 1.04</i>	<i>1.03, 1.06</i>	1.13, 1.04	1.08, 1.04	1.06, 1.04	<i>1.16, 1.04</i>
SMS	Repeated contact	1.03, 1.02	1.03, 1.02	0.99, 1.06	1.04, 1.02	0.99, 1.03	1.00, 1.03	0.99, 1.03
	Returned contact	0.99, 1.02	1.03, 1.00	1.00, 1.03	1.00, 1.02	0.98, 1.02	1.02, 1.01	0.99, 1.02
	Noncausal chain	1.02, 0.97	1.09, 0.95	<i>0.97, 1.02</i>	1.16, 0.97	1.04, 0.97	1.09, 0.96	<i>1.24, 0.97</i>
	Causal chain	0.99, 0.98	1.05, 0.97	0.95, 1.05	1.09, 0.98	1.00, 0.98	1.04, 0.97	<i>1.11, 0.98</i>
	Out-star	1.10, 1.00	1.17, 0.95	1.04, 0.98	1.35, 1.00	1.17, 1.00	1.21, 0.97	1.49, 1.01
	In-star	1.02, 1.00	1.13, 0.97	1.02, 0.98	1.18, 0.99	1.06, 0.99	1.16, 0.97	1.33, 0.99

The columns correspond to different attributes: age, sex, and payment type. The first value in each cell is the mean $r(m)$ for motifs where all nodes have the same attribute value (for example, all have the same age in column A). The second value gives the mean for all other motifs. If the first value is larger than the second, the motif has homophily with respect to those attributes: Cases where all nodes have the same value are relatively more common than others. Welch's t test was used to test for equality; bold denotes $P < 0.01$, and italic denotes $P < 0.05$ (including a Bonferroni correction corresponding to the number of tests in each table).

Table S2. Homophily for calls (upper) and SMS (lower) when either sex or payment type is fixed

	G = Fe	G = Ma	P = Po	P = Pr
Repeated contact	1.11, 1.11	1.13, 1.10	0.90, 1.18	1.27, 1.05
Returned contact	1.02, 1.01	1.02, 1.02	1.00, 1.02	0.95, 1.04
Noncausal chain	1.08, 1.02	1.01, 1.04	1.06, 1.00	0.96, 1.04
Causal chain	1.08, 1.01	0.98, 1.03	1.06, 0.97	0.92, 1.03
Out-star	1.10, 1.03	1.01, 1.04	<i>1.05, 1.03</i>	1.11, 1.03
In-star	1.11, 1.03	<i>1.01, 1.05</i>	<i>1.03, 1.06</i>	1.03, 1.05
Repeated contact	1.02, 1.03	1.04, 1.02	0.98, 1.04	1.01, 1.03
Returned contact	1.01, 1.02	1.05, 1.00	0.97, 1.03	1.04, 1.01
Noncausal chain	1.08, 0.96	<i>1.09, 0.97</i>	0.91, 1.03	1.04, 0.96
Causal chain	1.04, 0.98	<i>1.06, 0.98</i>	0.89, 1.05	1.05, 0.96
Out-star	1.18, 0.98	1.15, 1.00	1.05, 0.99	1.01, 1.01
In-star	1.09, 0.99	<i>1.19, 0.99</i>	1.02, 1.00	1.01, 1.00

The first value is average $r(m)$ for motifs where all nodes have the same value of sex (either female or male), or the same value of payment type (postpaid or prepaid). The second value is average $r(m)$ for all other motifs.