

Surges of Collective Human Activity Emerge from Simple Pairwise Correlations

Christopher W. Lynn,¹ Lia Papadopoulos,¹ Daniel D. Lee,² and Danielle S. Bassett^{1,3,4,5,6}

¹*Department of Physics and Astronomy, University of Pennsylvania,
Philadelphia, Pennsylvania 19104, USA*

²*Department of Electrical and Computer Engineering, Cornell Tech, New York, New York 10044, USA*

³*Department of Electrical and Systems Engineering, University of Pennsylvania,
Philadelphia, Pennsylvania 19104, USA*

⁴*Department of Bioengineering, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*

⁵*Department of Neurology, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*

⁶*Department of Psychiatry, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*



(Received 27 July 2018; revised manuscript received 11 December 2018; published 4 February 2019; corrected 7 March 2019)

Human populations exhibit complex behaviors—characterized by long-range correlations and surges in activity—across a range of social, political, and technological contexts. Yet it remains unclear where these collective behaviors come from or if there even exists a set of unifying principles. Indeed, existing explanations typically rely on context-specific mechanisms, such as traffic jams driven by work schedules or spikes in online traffic induced by significant events. However, analogies with statistical mechanics suggest a more general mechanism: that collective patterns can emerge organically from fine-scale interactions within a population. Here, across four different modes of human activity, we show that the simplest correlations in a population—those between pairs of individuals—can yield accurate quantitative predictions for the large-scale behavior of the entire population. To quantify the minimal consequences of pairwise correlations, we employ the principle of maximum entropy, making our description equivalent to an Ising model whose interactions and external fields are notably calculated from past observations of population activity. In addition to providing accurate quantitative predictions, we show that the topology of learned Ising interactions resembles the network of interhuman communication within a population. Together, these results demonstrate that fine-scale correlations can be used to predict large-scale social behaviors, a perspective that has critical implications for modeling and resource allocation in human populations.

DOI: [10.1103/PhysRevX.9.011022](https://doi.org/10.1103/PhysRevX.9.011022)

Subject Areas: Biological Physics, Complex Systems,
Statistical Physics

I. INTRODUCTION

In the study of human behavior, significant effort has focused on understanding the actions of one or two individuals at a time. It has been observed, for instance, that people engage in “bursts” of actions in quick succession [1–3], and significant effort has concentrated on understanding the correlated activity of pairs and triplets of individuals [3,4]. But if we broaden our perspective to an entire population, it becomes increasingly clear that humans also exhibit large-scale patterns of correlated activity. For example, urban transportation systems undergo surges of correlated activity known as traffic jams [5], first responders are required to handle correlated spikes in demand for emergency services [6], and internet and

telephone networks must be designed to withstand surges of collective activity [7,8]. But where do these large-scale patterns come from? Does it even make sense to discuss such distinct phenomena in the same breath?

Existing explanations for collective human behaviors have focused primarily on external mechanisms, such as fluctuations in urban traffic based on the time of the week [5] or spikes in demand for emergency services in response to natural disasters [6]. While external influences are an important part of the story, such explanations are inherently limited by their reliance on context-specific mechanisms like daily or weekly rhythms and natural disasters. By contrast, interactions between individuals are present in almost every human context, providing the possibility for a much more general explanation for the emergence of large-scale correlations. It is precisely this line of reasoning that has fostered vibrant efforts linking the study of social systems to tools and intuitions from statistical physics [9]. By adapting established models of collective behavior in physical systems, such as the Ising model and similar agent-based models, scientists have gained a deeper understanding of

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

the nature of collective behaviors in social systems. This program, for example, has resulted in Ising-like models of social dynamics and human cooperation [10–12], viral models aiding in the design of vaccination strategies [13], descriptions of the evolution of social networks [14], and statistical models of criminal activity [15,16].

Here, we draw inspiration from these seminal results to investigate the role of fine-scale correlations in generating large-scale patterns of human activity. Focusing on four data sets of human activity, from email and private message correspondence to physical contact and music streaming, we find that each population exhibits periods of intense collective activity, which cannot be explained by commonly used models that assume independence in human behavior [17–20]. Intuitively, these surges in activity could be driven by a common external influence, such as people’s daily and weekly schedules, or they could be driven by local correlations between pairs of individuals. To address the validity of the latter explanation, we quantify the collective impact of pairwise correlations by constructing a pairwise maximum entropy model that is formally equivalent to an Ising model from statistical mechanics. While the Ising model has previously been used to understand qualitative aspects of human activity [9,11,21–23], here, in order to make quantitative predictions, we calculate the specific external fields and pairwise interactions that best describe each population. In what follows, we show that this maximum entropy model (i) accurately predicts the frequencies of different patterns of collective human activity and (ii) bears a close resemblance to the network of interhuman communication within a population. Taken together, these results constitute an important step in the development of quantitative models of collective human behavior based on fine-scale correlations within a population. Such models, in turn, have important implications for resource allocation in communication [8] and transportation [5] networks, understanding social organization [24], and preventing viral epidemics [25].

II. NETWORK EFFECTS OF CORRELATIONS

As a salient example of collective human activity, we begin by studying patterns of email correspondence, focusing specifically on the email activity of 100 scientists at a European research institution over 526 days [26,27]. To understand the role of correlations in the timing of people’s actions—and in order to compare against other types of activities that are not directed from one individual to another [5–7,28,29]—we initially focus on the timing of sent emails while blinding our analysis to the email recipients. Importantly, this choice will later allow us to compare the architecture of functional interactions derived from our maximum entropy model with the network of communication within the population.

In a sufficiently small window of time Δt , each action appears to be binary—either individual i sent an email

($\sigma_i = 1$) or they were silent ($\sigma_i = 0$). By discretizing human activity in this way, we can begin to quantify correlations between people’s actions. We want the time window Δt to be as large as possible (to detect correlations between individuals) without being so large that individuals perform multiple actions within the same window. We find that nearly 90% of consecutive emails from the same individual are sent with at least two minutes in between [Fig. 1(a)], defining a natural timescale that we use as our Δt . Discretizing the data, as shown in Fig. 1(b), we produce a set of about 3.8×10^5 binary vectors σ , each of which captures the activity of the entire population within a given two-minute window.

The simplest and most common models of human activity assume that each individual behaves independently, implying that the number of people performing an action in

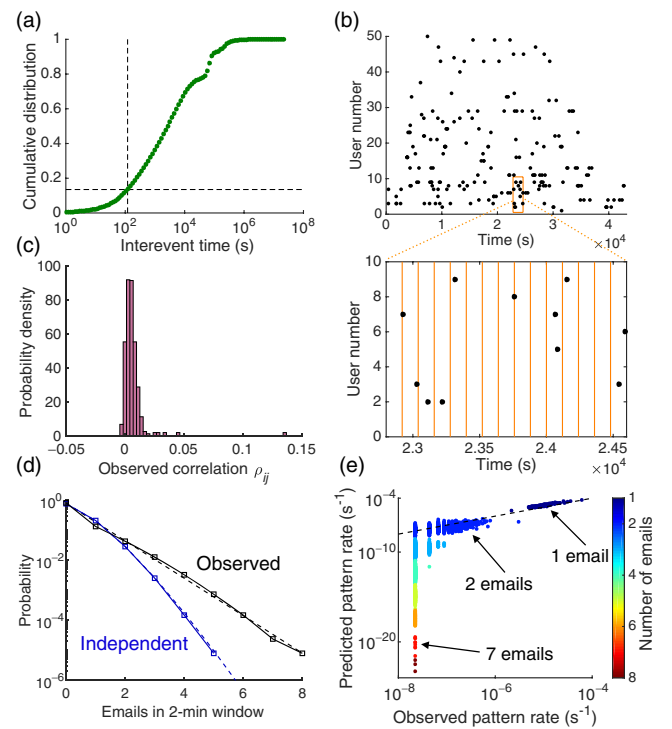


FIG. 1. Surges of human activity and failure of the independent approximation. (a) Distribution of interevent times for individuals in a network of email correspondence. The dashed lines indicate the proportion of interevent times less than two minutes. (b) Top: Activity of the 50 most active individuals over a half-day period, where each dot represents a sent email. Bottom: Network activity discretized into two-minute windows. (c) Histogram of Pearson correlation coefficients ρ_{ij} between activity time series for all pairs in the 100-person population. (d) Distribution of the number of emails sent in a given two-minute window (black) and the distribution after shuffling each person’s activity to eliminate correlations (blue). The dashed lines show an exponential distribution fit to the observed data (black) and a Poisson distribution fit to the shuffled data (blue). (e) The rate of each observed activity pattern, plotted against the approximate pattern rate assuming independent people. The dashed line indicates equality.

a given window follows a Poisson distribution [17]. Indeed, the Poisson distribution has been widely used to quantify the effects of various human actions, including telephone calls to a call center [18], internet activity [19], industrial accidents [17,18], and highway traffic flow [20]. In our population of email users, most pairs of individuals are only weakly correlated [Fig. 1(c)], suggesting that small groups should be well approximated by an independent model. However, if we extend the independent approximation to the entire population of 100 email users, it fails dramatically. While the Poisson distribution predicts a super-exponential dropoff in the number of actions performed in a given window, we find instead that human activity actually follows an exponential distribution [Fig. 1(d)]. This exponential distribution is characterized by a heavy tail, representing moments in time when many more people are sending emails than would be expected if they were behaving independently. Additionally, we report similar heavy-tailed distributions in separate data sets of private messages, physical contacts, and music streams (Figs. 10–12). For comparison, after shuffling the timing of the emails to eliminate correlations [30], we do not witness a window involving six or more active users [Fig. 1(d)], while we do observe about 1500 such instances in the original data set—nearly three per day.

The independent approximation also makes straightforward predictions for the rate of each activity pattern. Denoting the probability of individual i sending an email in a given two-minute window by $p_i(\sigma_i)$, the probability of observing a given activity pattern σ is simply predicted to be $P_1(\sigma) = \prod_i p_i(\sigma_i)$. This independent model severely underpredicts activity patterns involving three or more email users [Fig. 1(e)], and we find a similar discrepancy in a network of private messages [Fig. 10(c)]. In fact, under the independent model, each pattern of email activity involving seven active users should have only appeared roughly once every 10^{20} seconds—longer than the age of the Universe. We conclude that the independent approximation fails to explain the heavy-tailed nature of human behavior, characterized by surges of collective activity [5–8]. But where do these surges come from?

III. MAXIMUM ENTROPY MODEL OF HUMAN ACTIVITY

To improve upon the independent model, we must take into account correlations between individuals. Intuitively, such correlations could be driven by external influences such as daily and weekly rhythms [Fig. 2(a)], a hypothesis that has dominated existing explanations of large-scale human behaviors [5–8]. Alternatively, fine-scale correlations involving only a few individuals could build upon one another to have a strong impact on the population as a whole [Fig. 2(b)]. Here, we focus on the simplest possible correlations within a population—those between pairs of individuals—and ask whether these pairwise correlations can give rise to the large-scale patterns of activity that

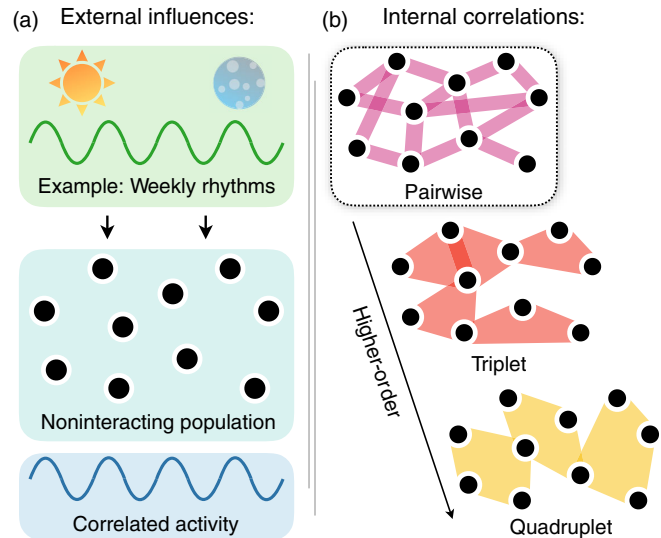


FIG. 2. External influences versus internal correlations. (a) An external mechanism—here taken to be weekly rhythms— influencing the activity of a population of noninteracting humans. Intuitively, circadian and weekly rhythms might influence people to send emails more frequently during the daytime and on weekdays, thereby inducing population-wide correlations. (b) Alternatively, population-wide correlations could arise from fine-scale interactions between individuals within a population. The set of all correlations forms a hierarchy, beginning with simple pairwise correlations between two individuals, followed by more complicated higher-order correlations involving three (triplet), four (quadruplet), or more individuals.

we observe in the data. As we will see, focusing on pairwise correlations represents a natural first step towards understanding emergent collective human activity, opening the door for straightforward generalizations to more complex higher-order correlations [Fig. 2(b)] [31,32].

We require a model that incorporates the observed pairwise correlations in the data while including as little information as possible about higher-order correlations between three, four, or more individuals. While it is not immediately obvious how one would construct such a model, Jaynes famously showed that an elegant solution lies in the principle of maximum entropy [33]: Among the infinite set of distributions consistent with a given set of correlations, the unique one that assumes as little information as possible about additional correlations is precisely the distribution with maximum entropy. This maximum entropy principle lies at the heart of equilibrium statistical mechanics [33,34] and has become increasingly popular as a tool for studying emergent phenomena in a range of complex systems, including networks of neurons in the brain [30,31], flocks of birds [35], protein structures [36], and gene coexpression patterns [37]. Despite this widespread adoption in biophysics, to our knowledge a similar data-driven approach has not been attempted previously in the social sciences.

Here, we consider the pairwise maximum entropy model, defined by the Boltzmann distribution

$$P_2(\boldsymbol{\sigma}) = \frac{1}{Z} \exp\left(\sum_i h_i \sigma_i + \frac{1}{2} \sum_{i \neq j} J_{ij} \sigma_i \sigma_j\right), \quad (1)$$

where the external fields h_i and pairwise interactions J_{ij} are Lagrange multipliers that ensure the model matches the observed individual activity rates and pairwise correlations in the data, respectively, and Z is the normalizing partition function. If we switch notation to $\sigma_i = \pm 1$, where $+1$ stands for activity and -1 for inactivity, P_2 is equivalent to the Ising model, which has long been used to simulate human dynamics in social networks [9,11,21–23]. However, while existing applications of the Ising model to human populations are based on metaphors about how people interact [11,12,21,23,38], we emphasize that our use of the Ising model is quantitatively rigorous in the sense that the external fields h_i and interactions J_{ij} are calculated to fit the observed activity of a given population (see Appendix D).

IV. MINIMAL CONSEQUENCES OF PAIRWISE CORRELATIONS

Calculations in the Ising model typically require summing over all 2^N activity patterns, where N is the number of elements in a system, prohibiting applications to large populations. Thus, it is common to construct a picture of the whole population by studying many different subpopulations [30], such as the 10 email users in Fig. 3(a). To quantify the explanatory power of pairwise correlations, we need meaningful ways to compare the accuracy of the maximum entropy model P_2 to that of the independent model P_1 . Toward this end, we use the Jensen-Shannon divergence $D_{JS}(Q||P)$ as a measure of distance from each of the model distributions (which we call Q) to the observed activity distribution P . Put simply, the Jensen-Shannon divergence represents the inverse of the number of independent samples needed to distinguish each model Q from the observed data [39]. Across 300 random groups of 10 users, we find that, on average, one would require 3.13×10^4 independent samples—over 43 days worth of data—to distinguish the pairwise model P_2 from the true distribution P [Fig. 3(b)]. By contrast, one would typically require 5 times fewer samples to distinguish the independent model P_1 from the observed data. Moreover, we find qualitatively similar results for individuals engaged in private messaging [Fig. 10(e)], face-to-face interactions [Fig. 11(c)], and online music streaming [Fig. 12(c)]. These observations suggest that the pairwise model provides a marked improvement in accuracy over the independent model.

We also wish to compare against a model representing the hypothesis that patterns of human activity are driven by external influences. While there are many external factors

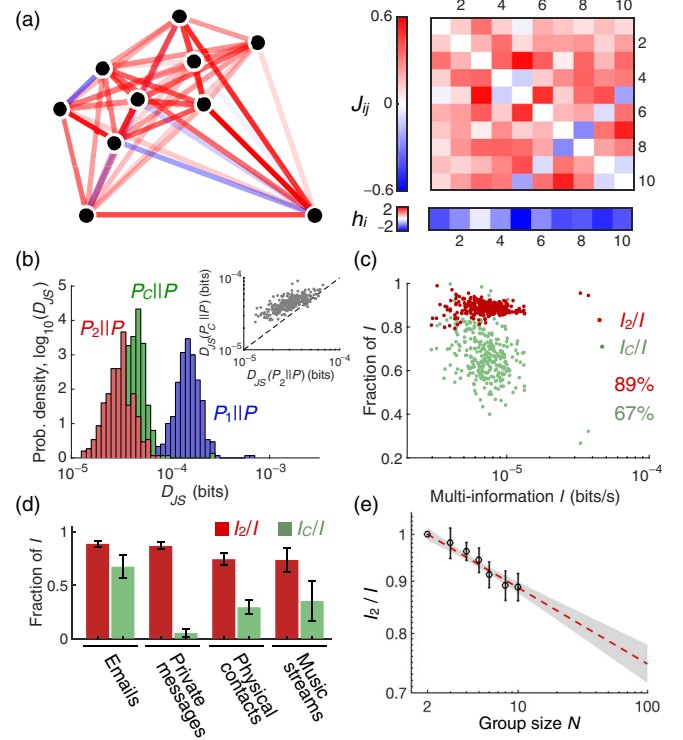


FIG. 3. The pairwise maximum entropy model accurately describes human behavior. (a) Learned Ising interactions J_{ij} and external fields h_i describing a random 10-person group in the email network. (b) Jensen-Shannon divergences between the true distribution P and the independent P_1 (blue), maximum entropy P_2 (red), and conditionally independent P_C (green) models. Histograms reflect estimates from 300 random groups of 10 individuals. Inset: $D_{JS}(P_2||P)$ versus $D_{JS}(P_C||P)$ for the 300 groups. The dashed line indicates equality. (c) Fraction of the network correlation (quantified by the multi-information I) captured by the maximum entropy (red) and conditionally independent (green) models, plotted against I for each group of 10 people. The multi-information is divided by Δt to remove dependence on the window size. (d) Fraction of the total correlation captured by the pairwise (red) and conditionally independent (green) models in four different modes of human activity: email correspondence, private messaging, physical interactions, and online music streaming. Error bars represent standard deviations over 300 random 10-person groups for the email and private message data sets, and over 200 groups for the physical contact and music streaming data sets. (e) Fraction of the multi-information in the email data captured by the maximum entropy model versus group size, where each data point is averaged over 300 randomly selected groups. The dashed line represents the best log-linear fit, with a 95% confidence interval indicated by the shaded region.

influencing human actions on a daily basis, from weather patterns to shifting demands at work, here we consider the most intuitive and well-studied external influence, namely, the impact of daily and weekly routines [see Fig. 2(a)] [5,7,8,40]. To formalize the hypothesis that activity patterns are driven by daily and weekly schedules, we consider the conditionally independent model P_C , wherein each

individual performs actions independently from all other individuals, but their activity rates are allowed to vary according to the time of the week [30,41] (see Appendix E). Compared to the conditionally independent model P_C , we find that the maximum entropy model P_2 is closer to the observed data (i.e., has a smaller Jensen-Shannon divergence from P) across 291 of the 300 groups [Fig. 3(c), inset]. This result is particularly notable when considering that P_2 only has 55 parameters for each group of 10 individuals, while P_C requires knowledge of each individual's email rate at each time during the week, totaling over 5×10^4 parameters.

The pairwise model accurately predicts the rates of particular activity patterns, but does it explain a majority of the total correlation in the population? To answer this question, we note that the total amount of correlation in the network, contributed by correlations between groups of users of all sizes, is quantified by the multi-information $I = S_1 - S$, where S_1 is the entropy of the independent distribution P_1 and S is the entropy of the observed distribution P [34] (see Appendix F). To determine the amount of multi-information that is contributed by pairwise correlations, it is useful to review the properties of maximum entropy models. For a population of N elements, we can define a sequence of maximum entropy models P_k that are consistent with all correlations up to the k th order, where $k = 1, 2, \dots, N$. These models form a hierarchy, from P_1 , in which all elements are independent, up to P_N , which is an exact description of the observed activity. As we scale this hierarchy, the entropies S_k of the distributions decrease monotonically toward the true entropy ($S_1 \geq S_2 \geq \dots \geq S_N = S$), and the combined contribution of all k th-order correlations is quantified by the entropy difference $I_k = S_{k-1} - S_k$. We note, for instance, that these entropy differences sum to the full multi-information: $I_2 + \dots + I_N = I$. Thus, the problem of determining how much of the total correlation in the data stems from pairwise correlations formally reduces to calculating the proportion of the multi-information I that is accounted for by the reduction in entropy from pairwise correlations (i.e., $I_2 = S_1 - S_2$).

We observe that pairwise correlations account for a striking $I_2/I \approx 89\%$ of the total correlation in groups of 10 users [Fig. 3(c)]. In turn, this observation implies that the contributions of all other higher-order correlations, $I_3 + \dots + I_N$, only combine to account for the remaining 11% of the multi-information. Meanwhile, the amount of correlation attributable to daily and weekly rhythms is represented by the entropy difference $I_C = S_1 - S_C$, where S_C is the entropy of the conditionally independent model P_C . This popular explanation for collective human behavior is consistently less effective than the maximum entropy model at capturing the correlations in the data [$I_C/I \approx 67\%$; Fig. 3(c)]. Importantly, we show (i) that these results are robust to both reasonable variation in the time

window Δt used to discretize the data [Appendix B 1, Fig. 7] and differences in the set of individuals selected for analysis [Appendix B 2, Fig. 8] and (ii) that the maximum entropy model is relatively consistent over time [Appendix B 3, Fig. 9]. Moreover, we verify that similar results hold in separate data sets of private messages [Appendix C 1, Fig. 10], physical contacts between individuals [Appendix C 2, Fig. 11], and music streaming online [Appendix C 3, Fig. 12], as summarized in Fig. 3(d). In the data set of private messages, for instance, the pairwise model captures nearly the same amount of correlation as in the population of email users ($I_2/I \approx 87\%$), while people's daily and weekly rhythms explain very little of the correlation [$I_C/I \approx 5\%$; Fig. 3(e)]. Interestingly, this difference in I/I_C between email activity and private messages [Fig. 3(c)] reflects the commonly-held intuition that email activity is moderately tied to people's work and leisure schedules, while private messages are not.

We are ultimately interested in understanding the role of pairwise correlations in driving large-scale surges of activity in the entire 100-person population. With this goal in mind, we calculate the fraction I_2/I in groups of email users increasing in size from $N = 2$ through 10. For small groups and relatively weak correlations, as the group size increases, we expect the multi-information I to increase in proportion to the entropy difference I_2 [30]. Indeed, we find that the fraction I_2/I remains nearly constant as the groups grow in size ($I_2/I \propto N^{-0.075 \pm 0.005}$). Extrapolating to the entire 100-person population, we find with 95% confidence that pairwise correlations account for 72%–78% of the total multi-information in the data [Fig. 3(d)]. This fraction is especially large when considering the exponential number of possible higher-order correlations ($\sim 2^N$) for populations of increasing size N . We conclude that large-scale patterns of behavior, across several distinct modes of human activity, can be robustly understood as emerging from an underlying network of pairwise correlations.

V. MODELING AN ENTIRE POPULATION

Our analysis of relatively small groups indicates that the pairwise maximum entropy model can capture a majority of the correlation structure in groups of up to 100 individuals. This result, in turn, suggests that the heavy-tailed nature of collective human behavior [Fig. 1(d)]—characterized by surges of activity—might emerge organically from pairwise correlations. To test this prediction directly, we must extend the pairwise maximum entropy model to include the entire population of 100 email users. In order to learn the appropriate Ising interactions J_{ij} and external fields h_i for all 100 people, we leverage recent advances in stochastic gradient descent from statistical physics [42] and machine learning [43], avoiding the exponential complexity of standard Ising

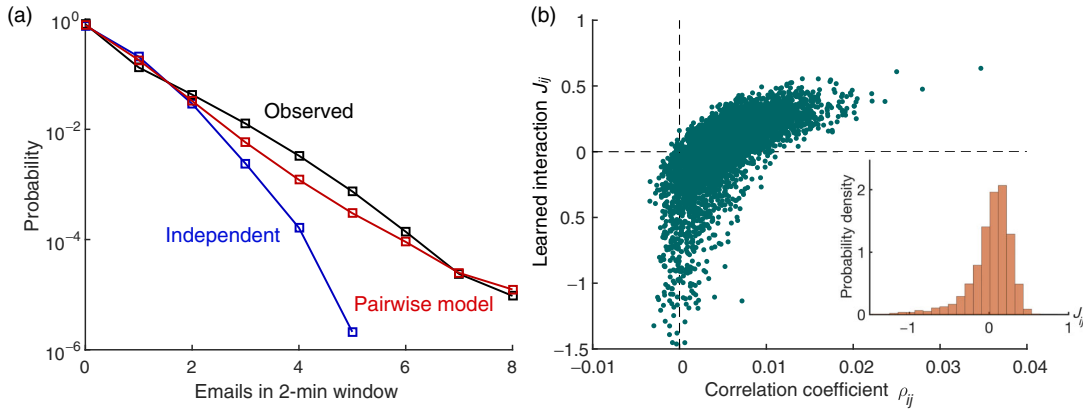


FIG. 4. Surges of collective activity are captured by pairwise correlations. (a) Distribution of the observed number of emails in a given two-minute window (black), the prediction of the independent model (blue), and the prediction of the pairwise maximum entropy model (red). (b) Scatter plot illustrating the relationship between the observed pairwise correlations in the data ρ_{ij} and the learned Ising interactions J_{ij} for all pairs in the 100-person population. Inset: Histogram of the learned interactions.

calculations [see Appendix D; Fig. 13]. Figure 4(a) shows that the pairwise model successfully captures the heavy-tailed nature of human activity, accurately predicting the frequencies of activity surges involving up to seven and eight individuals.

To understand how a network of simple pairwise correlations can generate large-scale spikes in activity, it is useful to study the structure of the Ising parameters in the maximum entropy model [Eq. (1)]. We note that each external field h_i either biases individual i toward activity ($h_i > 0$) or toward inactivity ($h_i < 0$). Meanwhile, each Ising interaction J_{ij} either influences individuals i and j to perform actions at the same time ($J_{ij} > 0$) or at different times ($J_{ij} < 0$). Here, we draw an important distinction between the learned interactions J_{ij} in the maximum entropy model and the observed pairwise correlations ρ_{ij} in the data: While each pairwise correlation quantifies the frequency with which two individuals perform actions at the same time, each Ising interaction represents a functional influence between two individuals to synchronize their activity, thereby inducing a pairwise correlation. Interestingly, while correlations in the network are weak and almost exclusively positive [Fig. 1(c)], the Ising interactions maintain a large amount of heterogeneity [Fig. 4(b), inset], with almost an equal number of positive and negative interactions. Indeed, the learned pairwise interactions depend highly nontrivially on the corresponding pairwise correlations in the data [Fig. 4(b)]. Importantly, the presence of competing positive and negative interactions generates “frustration,” as in spin glasses [44], wherein triplets of individuals cannot find a combination of activity and inactivity that simultaneously satisfies all of their interactions. This frustration gives rise to a complex energy landscape of activity patterns with many different local minima, some of which correspond to patterns involving many more active individuals than would be expected under the independent model, thus giving rise to the

heavy-tailed behavior in Fig. 4(a). Intriguingly, such frustrated interactions have previously been hypothesized to drive a number of social phenomena [9], such as the formation of coalitions [45]. By calculating the specific Ising parameters that describe each population, and by identifying the presence of competing positive and negative interactions [Fig. 4(b), inset], our work provides rigorous evidence for these long-standing hypotheses.

VI. ROLE OF INTERHUMAN COMMUNICATION

Thus far, we have focused on understanding correlations in the timing of actions, without knowledge of who each person is interacting with in the population. Fundamentally, the Ising interactions J_{ij} are merely learned parameters that ensure consistency with the observed pairwise correlations in the network. However, it is tempting to imbue them with physical significance, interpreting these functional interactions as comprising a network of real-world influences between individuals. For previous applications of maximum entropy models in neuroscience [30,31] and biology [35–37], because comparisons with ground truth interactions are often infeasible, any physical meaning attributed to the learned interactions J_{ij} has remained, at its core, an analogy. By contrast, in the context of email activity, we automatically know a subset of the ground truth interactions—namely, the network of email communication between individuals. Although it is appealing to suspect that the learned Ising interactions are closely related to the structure of email correspondence in the data, we emphasize that this need not be the case. There is an array of circumstances that could influence the activity of two individuals to become correlated, from common functional roles in the network to shared communication with an external third party. Furthermore, even if correlations do arise from direct communication, this communication could take on many forms that do not appear in the data

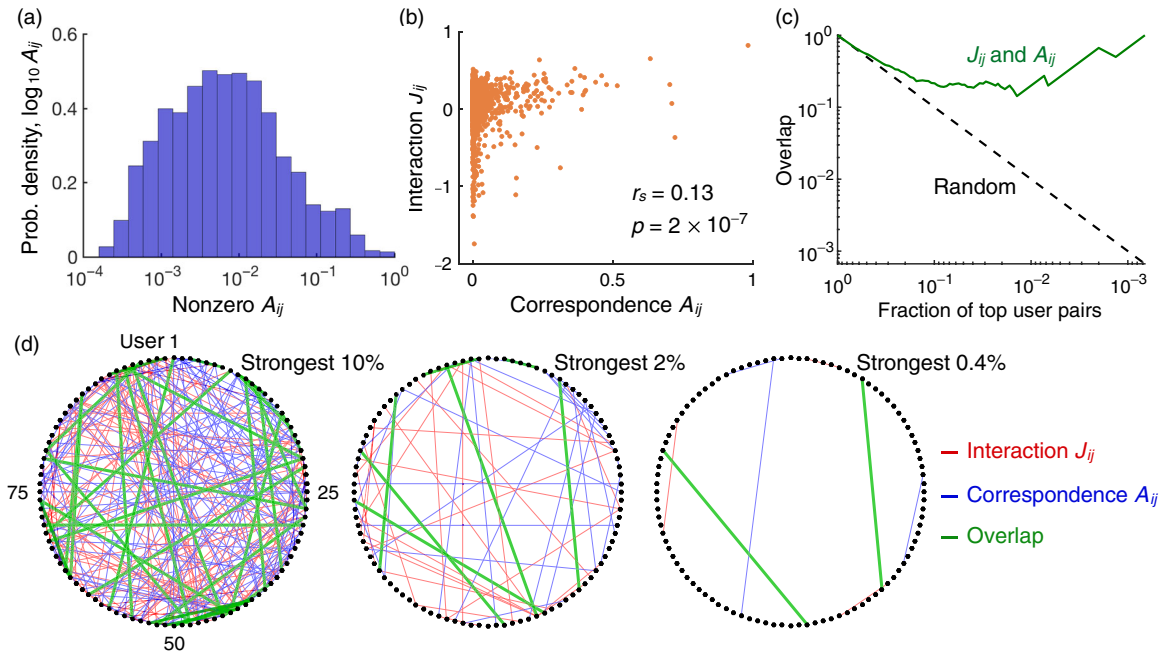


FIG. 5. The learned pairwise interactions uncover pathways of ground truth communication. (a) Histogram of correspondence rates A_{ij} between all pairs of individuals that exchanged at least one email. (b) Scatter plot of the learned Ising interactions versus email correspondence rates for pairs that exchanged at least one email. Importantly, J_{ij} and A_{ij} are significantly correlated with Spearman's correlation coefficient $r_s = 0.13$ ($p = 2 \times 10^{-7}$). (c) Overlap between the strongest interactions J_{ij} and most frequently corresponding pairs A_{ij} as a function of the fraction of pairs being considered. The dashed line indicates the overlap with a random selection of user pairs. (d) Structure of the strongest pairwise interactions (red), highest correspondence rates (blue), and overlap between the two (green) for all 100 individuals. The three networks represent the strongest 10% (left), 2% (middle), and 0.4% (right) of user pairs.

set, including face-to-face contact, texts, calls, or other online avenues.

Keeping in mind these reasons for guarded optimism, here we compare the learned interactions J_{ij} from our maximum entropy model with the network of email traffic between individuals. Letting $n_{i \rightarrow j}$ denote the number of emails sent from person i to person j , and letting $n_i = \sum_j n_{i \rightarrow j}$ denote the total number of emails sent by person i , we define the correspondence rate between two people i and j to be $A_{ij} = (n_{i \rightarrow j} + n_{j \rightarrow i}) / (n_i + n_j)$. In other words, A_{ij} represents the fraction of the $n_i + n_j$ emails sent by person i and person j that were addressed to each other. We find that most correspondence only accounts for around 1% of a pair's total email communication, while a small number of pairs communicate almost exclusively with one another [Fig. 5(a)]. Considering all pairs of people that exchanged at least one email ($A_{ij} > 0$), we find that the learned Ising interactions J_{ij} are significantly correlated with the correspondence rates A_{ij} in the data [Spearman's correlation coefficient $r_s = 0.13$, $p = 2 \times 10^{-7}$; Fig. 5(b)]. This relationship between the learned Ising interactions and the ground truth communication in the population is particularly interesting after reflecting on the myriad ways

in which these two networks could have remained unrelated, as described above.

To fully appreciate the strength of the relationship between J_{ij} and A_{ij} , we focus on the fraction f of the strongest pairwise interactions and correspondence rates in the population. These two thresholded networks overlap significantly [Fig. 5(c)], with the strongest 1% of Ising interactions exhibiting a 20% overlap with the top 1% of frequently communicating pairs—20 times higher than expected if J_{ij} and A_{ij} were independent. This overlap becomes even more pronounced as we increase the threshold [Fig. 5(d)], such that the single strongest maximum entropy interaction in the entire population corresponds precisely to the pair of individuals that communicate most frequently. This relationship between J_{ij} and A_{ij} provides a compelling mechanistic interpretation for the Ising interactions in our maximum entropy model; namely, frequent communication between a pair of individuals (quantified by A_{ij}) acts as an influence to synchronize their activity (quantified by J_{ij}). As demonstrated in previous sections, the resulting pairwise correlations, in turn, can generate the types of large-scale correlations and surges in human activity that are ubiquitous in the modern world [5–8,28,29].

VII. CONCLUSIONS AND FUTURE DIRECTIONS

Despite the widespread investigation of fine-scale correlations as the building blocks of large-scale behavior in complex systems throughout physics [33,34], neuroscience [30,31,46], and biology [35–37], a similar quantitative approach to human dynamics has been notably lacking. Here, we provide an important step toward the ultimate goal of understanding the role of fine-scale correlations in generating large-scale patterns of human activity. Studying four data sets that reflect the diversity of human activity, we first show that all populations exhibit surges of collective activity, a phenomenon that has become the subject of intense research focus [5–8,28,29]. Importantly, these surges in activity cannot be accounted for by commonly used models that assume independence in human behavior [17–20]. To understand where surges in activity come from, we consider the possibility that large-scale patterns arise naturally from combinations of simple pairwise correlations between individuals. To formalize this hypothesis, we utilize the principle of maximum entropy from information theory, deriving a pairwise maximum entropy model of human activity that is formally equivalent to an Ising model. Interestingly, this maximum entropy model accounts for 72%–78% of the total correlation in a 100-person population of email users [Fig. 3(e)] and accurately predicts the heavy-tailed distribution of activity surges [Fig. 4(a)]. Additionally, we demonstrate that the Ising interactions in our model closely resemble the network of interhuman communication within the population. This close relationship between functional interactions and ground truth communication suggests an intuitive mechanism driving pairwise correlations.

Just as emergent phenomena have garnered significant attention in the natural sciences [30–37], we anticipate that similar approaches will prove fruitful in the development of accurate models of large social systems. Importantly, while a majority of existing research has focused on the impacts that external influences have on human populations [5,6], these explanations are fundamentally limited by their reliance on context-specific mechanisms [7,8]. By contrast, interactions between humans are present in almost every context, and, as we have demonstrated, these interactions can build upon one another to have a large-scale impact on the behavior of an entire population. In this way, thinking carefully about the role of fine-scale correlations in activity can have quite general implications for resource allocation in communication [8] and transportation [5] networks, understanding social organization [24], and preventing viral epidemics [25].

To conclude, we point out a number of limitations of our analysis that highlight important directions for future work. First, we remark that, given the diversity of experiences that shape human actions, it would be naive to conclude that all collective behaviors only emerge from internal correlations. To the contrary, it has been well established that external

influences play an important role in predicting a number of collective human behaviors [5–8,28,29]. Therefore, future work should investigate the interplay between external influences and internal interactions in human populations. Such an investigation would likely benefit from advances in control theory and influence maximization [47,48], which have recently been used to predict the propagation of external influences in Ising networks [23,38,49]. Second, we note that our investigation has focused primarily on pairwise correlations. While these simplest correlations represent a logical first step, our results do not rule out the possibility that higher-order correlations could also have an important impact on large-scale behavior. Practically speaking, the primary difficulty in studying such higher-order correlations lies in determining which to include in a maximum entropy model, as there exist $\binom{N}{k}$ different choices for each k th-order correlation (a number that grows nearly exponentially with k). Fortunately, to handle this explosion of parameters, recent advances in neuroscience have produced tractable techniques for generating sparse higher-order maximum entropy models [31]. Such higher-order models represent systematic generalizations of the methods presented here and could prove vital for understanding the large-scale impacts of triplet and quadruplet correlations [Fig. 2(b)], which are thought to encode important organizational features in human populations [4] (see Appendix G for an extended discussion).

ACKNOWLEDGMENTS

D. S. B. and C. W. L. acknowledge support from the John D. and Catherine T. MacArthur Foundation, the Alfred P. Sloan Foundation, the ISI Foundation, the Paul Allen Foundation, the Army Research Laboratory (W911NF-10-2-0022), the Army Research Office (Bassett-W911NF-14-1-0679, Grafton-W911NF-16-1-0474, DCIST-W911NF-17-2-0181), the Office of Naval Research, the National Institute of Mental Health (2-R01-DC-009209-11, R01-MH112847, R01-MH107235, R21-MMH-106799), the National Institute of Child Health and Human Development (1R01HD086888-01), National Institute of Neurological Disorders and Stroke (R01 NS099348), and the National Science Foundation (BCS-1441502, BCS-1430087, NSF PHY-1554488, and BCS-1631550). L. P. is supported by the NSF. D. D. L. acknowledges support from the NSF, NIH, DOT, ARL, AFOSR, ONR, and DARPA. The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

APPENDIX A: DATA PREPROCESSING

Here, we discuss the details of how the email data are processed, noting that the other data sets follow in an analogous fashion. In total, the data set contains the email correspondence between 986 members of a European

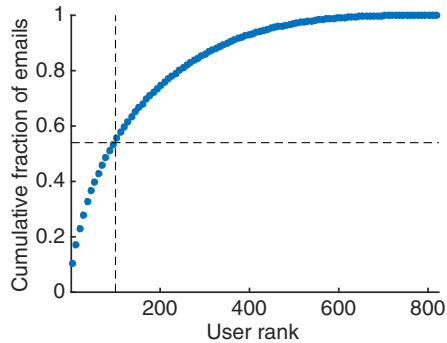


FIG. 6. Cumulative distribution of emails versus the activity rank of the users. The 100 most active individuals account for 56% of the emails in the network (dashed lines).

research institution over 526 days [26]. We focus on the 100 most active individuals, roughly corresponding to the members of the population that sent, on average, at least one email per day (Fig. 6). To quantify correlations

between different individuals, we must discretize the data into time bins of width Δt . To choose a suitable bin width, we notice that 90% of consecutive emails from the same individual are sent with at least two minutes in between [Fig. 1(a)], defining a natural timescale that we use as our Δt . Discretizing the 526-day data set into two-minute bins, we produce a set of $\sim 3.8 \times 10^5$ binary patterns $\{\sigma\}$ that define the behavior of our population.

APPENDIX B: ROBUSTNESS OF THE PAIRWISE MODEL

In Appendix A, we provided first-principles justifications for focusing on the 100 most active individuals in the email data set and for discretizing the data into bins of width $\Delta t = 2$ minutes. Here, we verify that the success of the pairwise maximum entropy model is robust to reasonable variations in these choices.

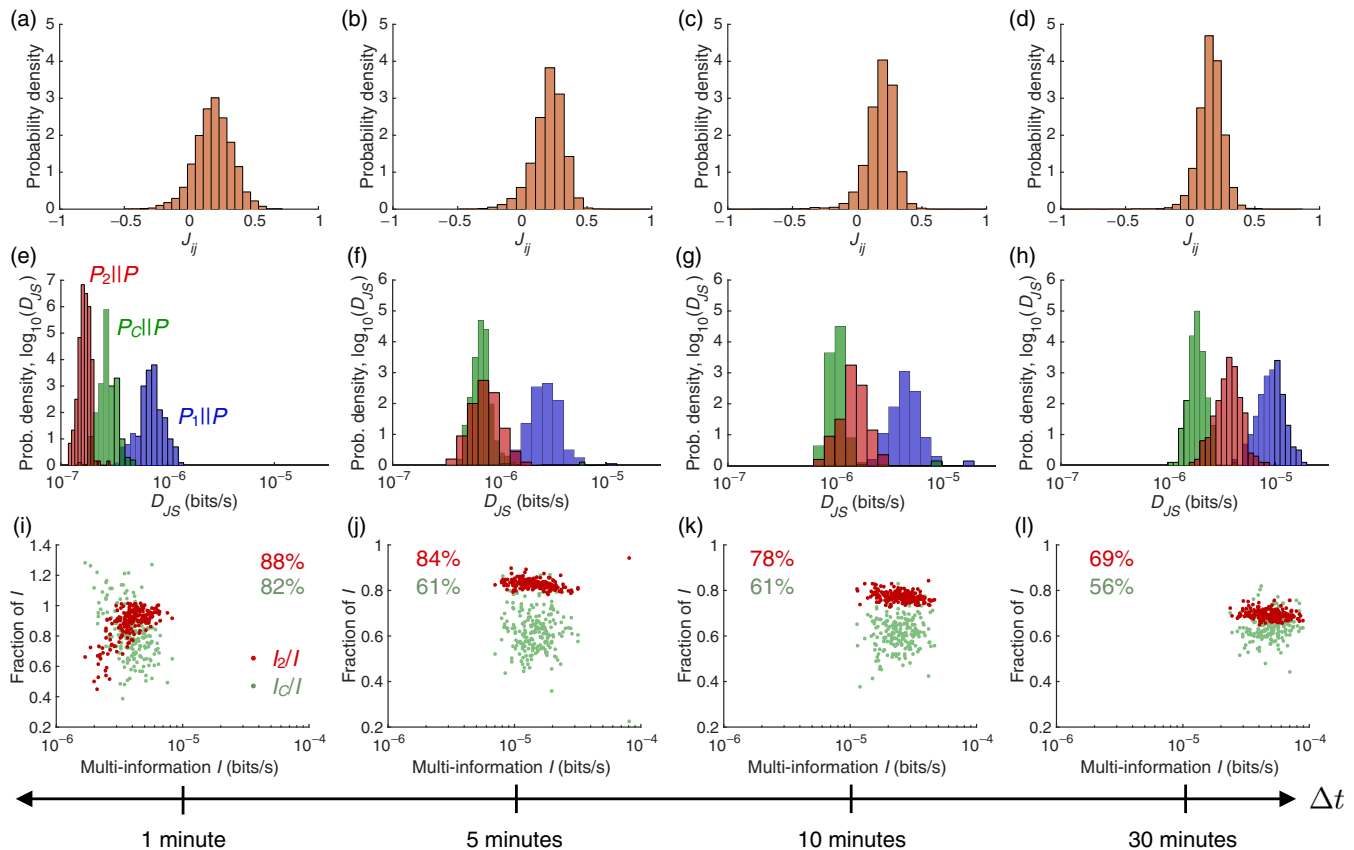


FIG. 7. Dependence of the pairwise maximum entropy model on the bin width. (a)–(d) Distributions of pairwise couplings for 200 different 10-person groups selected from the 100 most active individuals in the email data set. From left to right, the data are discretized into bins of width $\Delta t = 1, 5, 10,$ and 30 minutes. (e)–(h) Jensen-Shannon divergences between the observed distribution over activity patterns P and the independent P_1 (blue), maximum entropy P_2 (red), and conditionally independent P_C (green) models. The distributions are taken over the 200 groups from panels (a)–(d). (i)–(l) Fraction of the network correlation captured by the maximum entropy (red) and conditionally independent (green) models, plotted against the full network correlation, quantified by the multi-information I . The average percentage of the multi-information captured by each model is displayed in the upper corner. Each dot represents a different group of 10 people, and I is divided by Δt to remove dependence on the window size.

1. Dependence on the bin width

We investigate the dependence of the pairwise maximum entropy model on the bin width Δt used to discretize the email activity. Throughout, we focus on the 100 most active individuals, and we consider bin widths of $\Delta t = 1, 5, 10,$ and 30 minutes. For each value of Δt , we randomly select 200 different groups of 10 individuals and fit a pairwise maximum entropy model to describe each group. As Δt increases, we witness more windows involving multiple active individuals, thereby strengthening the correlations that we observe in the discretized data. In turn, these stronger correlations give rise to Ising interactions J_{ij} that are more positive and sharply peaked [Figs. 7(a)–7(d)]. In Figs. 7(e)–7(h), we show that the true distribution of activity is approximately 5 times closer to the maximum entropy model P_2 than to the independent model P_1 across

all values of Δt considered, demonstrating the consistency of the pairwise model in predicting human behavior. On the other hand, the performance of the conditionally independent model P_C increases significantly as Δt increases, even outperforming the pairwise model for $\Delta t \geq 10$ minutes. We note, however, that for such large bin widths, people often send multiple emails within the same window, and treating the data as binary may not be justified. In Figs. 7(i)–7(l), we see that the pairwise model captures nearly all of the multi-information in the 10-person groups across all choices for Δt . By contrast, the conditionally independent model consistently captures a smaller fraction of the multi-information in the data. Furthermore, for $\Delta t = 1$ minute, the conditionally independent model has lower entropy than the data themselves (i.e., $I_C/I > 1$) for 30 of the 200 groups, which is a clear indication that the model is overfitting the data.

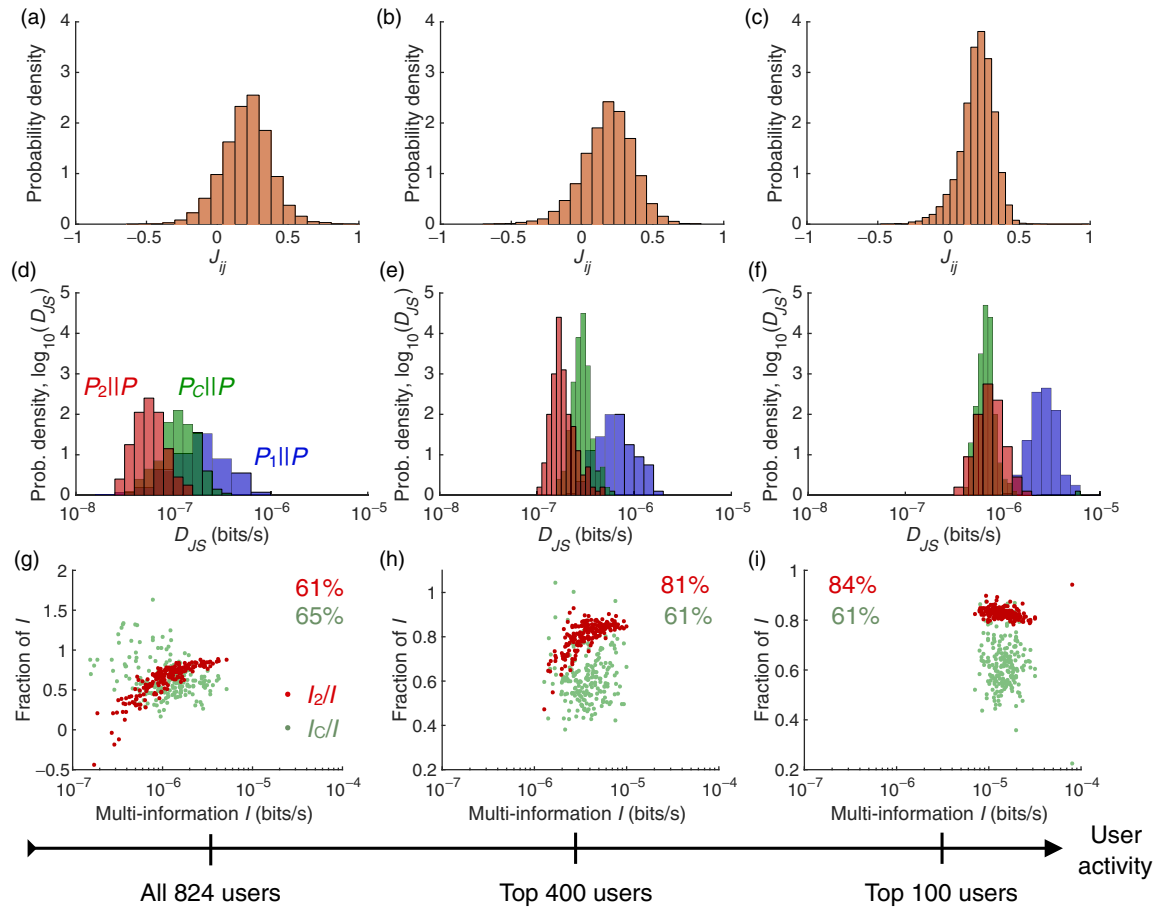


FIG. 8. Dependence of the pairwise model on the set of individuals chosen for analysis in the email data set. (a)–(c) Distributions of pairwise interactions for 200 different groups of 10 individuals, where the data are discretized with bin width $\Delta t = 5$ minutes. From left to right, the 200 groups are chosen from among all 824 people that sent at least one email, the 400 most active individuals, and the 100 most active individuals, respectively. (d)–(f) Jensen-Shannon divergences between the observed distribution over activity patterns P and the independent P_1 (blue), maximum entropy P_2 (red), and conditionally independent P_C (green) models. The distributions are taken over the 200 groups of users. (g)–(i) Fraction of the network correlation captured by the pairwise maximum entropy (red) and conditionally independent (green) models, plotted against the full network correlation, quantified by the multi-information I . The average percentage of the multi-information captured by each model is displayed in the upper corner. The multi-information is divided by Δt to remove dependence on the window size.

2. Dependence on the individuals being analyzed

We investigate the dependence of the maximum entropy model on the set of individuals chosen for analysis. In particular, we consider 200 different 10-person groups selected from among the 100 most active email users, the 400 most active users, and all 824 users that sent at least one email. Throughout this section, the bin width is fixed at $\Delta t = 5$ minutes. As we focus on more active individuals, the observed correlations become stronger, which is reflected in the fact that the distribution of learned interactions J_{ij} among the top 100 individuals is more sharply peaked and positive than the pairwise interactions between the top 400 and all 824 individuals [Figs. 8(a)–8(c)]. In Figs. 8(d)–8(f), we again find that the pairwise model is approximately 5 times closer to the true distribution than the independent model across all three subpopulations. By contrast, the conditionally independent model performs nearly as well as the pairwise model among the 100 most active individuals but provides only marginal improvements over the independent model for all 824 individuals. The failure of the conditionally independent model in describing the entire 824-person population is not surprising given that most individuals sent less than one email every five days, leaving daily and weekly rhythms with little to no predictive power.

We now consider the fraction of the multi-information captured by each model. For all 824 individuals, Fig. 8(g)

shows that the conditionally independent model captures a slightly larger fraction of the multi-information than the maximum entropy model; however, P_C erroneously includes more correlation than the data themselves (i.e., $I_C/I > 1$) for 20 of the 200 groups of 10 people, indicating that the model is overfitting the data. For both the top 100 and 400 most active individuals, the maximum entropy model captures a significantly larger fraction of the network correlation than the conditionally independent model [Figs. 8(h)–8(i)].

We conclude that the predictions of the pairwise maximum entropy model are robust to variations in both the bin width Δt and the set of individuals chosen for analysis.

3. Consistency of the pairwise model over time

By employing the pairwise maximum entropy model in Eq. (1), we implicitly assume that the population activity can be modeled as a stationary distribution, that is, that the local fields h_i and interactions J_{ij} do not change over time. Here, we test this assumption explicitly while noting that the development of time-evolving maximum entropy models is an important direction for future work (see Appendix G 3 for an extended discussion). Specifically, we wish to determine if the Ising parameters describing one portion of the email activity resemble those describing another portion of the activity. To do so, we divide the data set into two halves corresponding roughly to the first and last 263 days of email

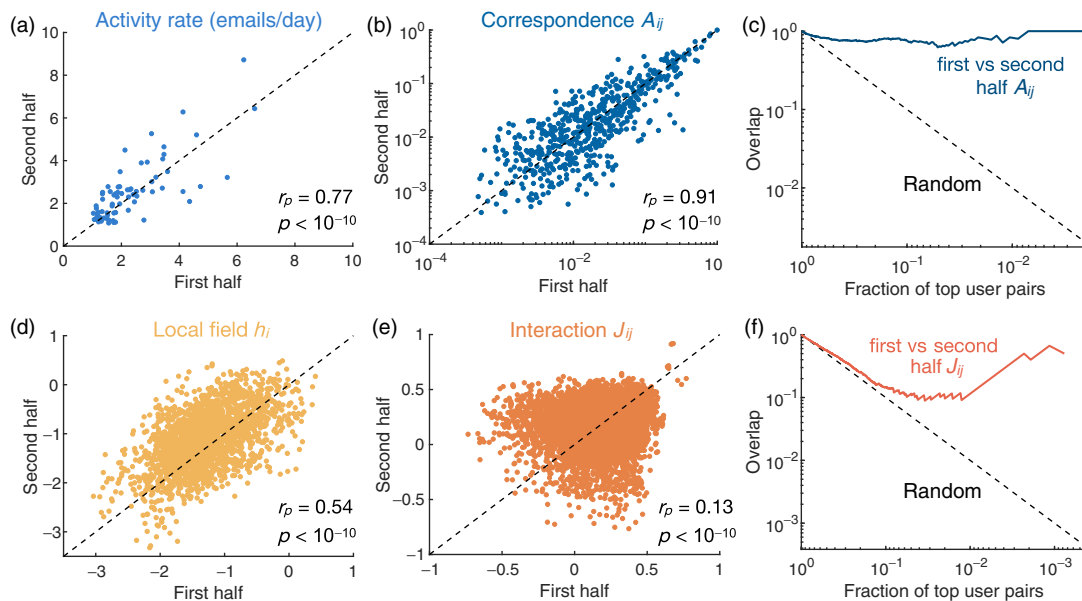


FIG. 9. Consistency of the pairwise maximum entropy model over time. (a) Comparison of email user activity rates in the first half versus the second half of the data set; the dashed line indicates equality. (b) Correspondence rates A_{ij} between pairs of users are strongly correlated across the two halves of the data set. (c) Overlap between the most frequently corresponding pairs of users in the first half and those in the second half as a function of the fraction of pairs being considered. The dashed line indicates the overlap with a random selection of user pairs. (d) For 200 random groups of 10 individuals, we compare the local fields h_i of pairwise maximum entropy models fit to either the first or second half of the email data. (e) For the same 200 random groups, we compare the Ising interactions J_{ij} of the pairwise models fit to the two halves of the data set. (f) For each half of the data set, we average the interactions J_{ij} over all 200 groups and plot the overlap between average interaction networks as a function of the fraction of user pairs being considered. As in panel (c), the dashed line indicates the overlap with a random selection of pairs.

activity. Figures 9(a)–9(c) show that the statistics describing the population activity remain remarkably consistent over time, with both the user activity rates and pair correspondence rates A_{ij} being strongly correlated between the two halves of data (Pearson’s correlations $r_p = 0.77$ for the activity rates and $r_p = 0.91$ for the correspondence rates).

To study the consistency of the maximum entropy model, we randomly select 200 different 10-person groups from among the 74 users that sent at least one email in both halves of the data set; we then learn pairwise models describing each group for each half of data. Figures 9(d) and 9(e) show that the local fields h_i and interactions J_{ij} modeling the population activity are significantly correlated over time (Pearson’s correlations $r_p = 0.54$ for the local fields and $r_p = 0.13$ for the interactions). The consistency of the Ising interactions J_{ij} between the two halves of data becomes even more apparent when we focus on the strongest interactions in the population [Fig. 9(f)]. Together, these results indicate that the patterns of population activity remain relatively consistent over time, justifying our application of the stationary maximum entropy model as a first step toward more complex dynamical models.

APPENDIX C: OTHER MODES OF HUMAN ACTIVITY

In the main text, our analysis focused primarily on a data set of email activity. Here, we independently verify the ability of the pairwise maximum entropy model to quantitatively describe collective human behavior in three other data sets representing a diverse range of human activities.

1. Private messages

We first consider a data set of around 6×10^5 private messages sent between 1899 students at U.C. Irvine over the span of 193 days [27]. As in the context of email activity, we focus on the individuals that sent, on average, at least one message per day, corresponding to the 66 most active students in the population. To choose an appropriate bin width, we consider the distribution of time gaps between consecutive messages from the same student [Fig. 10(a)]. Comparing against the equivalent distribution in the email data set [Fig. 6(b)], we notice that many more private messages than emails are sent with short gaps ($\lesssim 1$ minute) in between. This bursty behavior indicates that

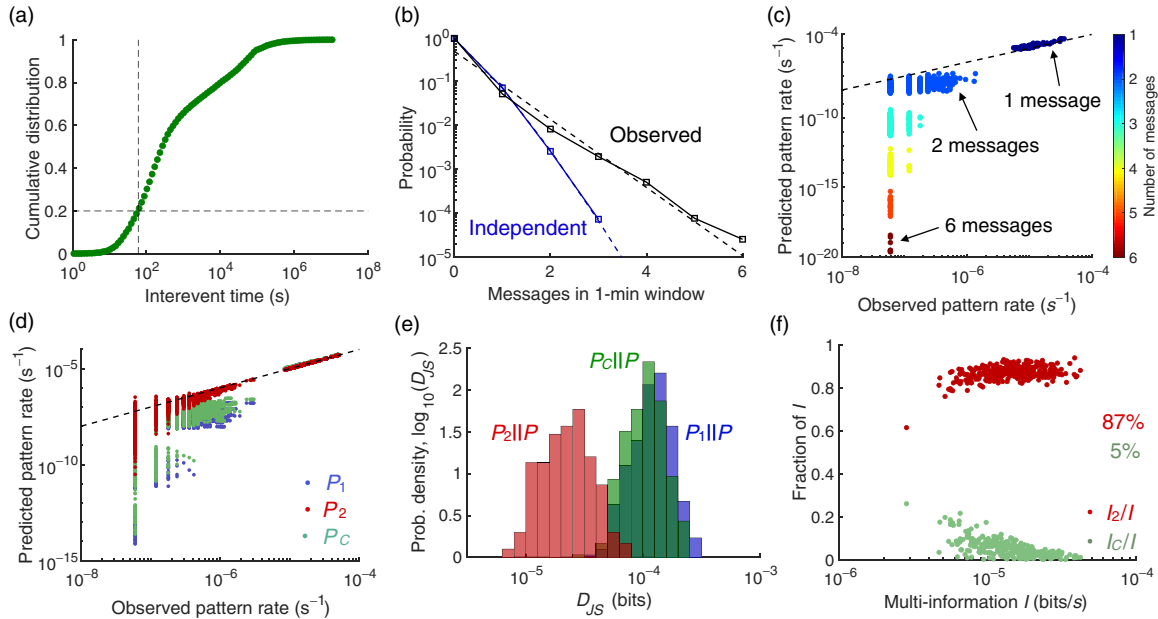


FIG. 10. Performance of the pairwise maximum entropy model in a data set of private messages. (a) Cumulative distribution of interevent times for the 66 most active individuals. Approximately 80% of consecutive messages from the same person are sent with at least one minute in between (dashed lines). (b) Distribution of the messages sent in a given one-minute window in the data set (black) and after shuffling individuals’ activities to eliminate correlations (blue). Dashed lines indicate an exponential fit to the observed data (black) and a Poisson fit to the shuffled data (blue). (c) The rate of each observed activity pattern, plotted against the approximate rate under the independent model P_1 . The dashed line indicates equality. (d) The rate of each observed activity pattern across 300 randomly selected groups of 10 individuals against the approximate rates under the independent model P_1 (blue), the pairwise maximum entropy model P_2 (red), and the conditionally independent model P_C (green). The dashed line indicates equality. (e) Jensen-Shannon divergences between the true distribution P and the independent P_1 (blue), maximum entropy P_2 (red), and conditionally independent P_C (green) models. The histograms reflect estimates from the 300 10-person groups. (f) Fraction of the network correlation (i.e., multi-information I) captured by the pairwise (red) and conditionally independent (green) models, plotted against the full multi-information. We note that I is divided by Δt to remove the dependence on window size.

the private messages serve as a more conversational communication medium than emails, a fact that will later help in understanding the impact of daily and weekly rhythms. Because of the bursty nature of private messages, we reduce our bin width to $\Delta t = 1$ minute, yielding a data set of about 2.8×10^5 binary activity patterns.

As in the network of email correspondence, the independent model P_1 fails to explain the collective behavior in the private message population [17–20]; while the independent model predicts a super-exponential dropoff in the number of active individuals in a given window, we find that the distribution of private messages is actually heavy tailed, fitting closely to an exponential distribution [Fig. 10(b)]. Additionally, in Fig. 10(c), we see that the independent model dramatically underpredicts patterns involving two or more active individuals. To improve upon the independent model, we again consider two competing hypotheses: (i) that large-scale patterns emerge from an aggregation of simple pairwise correlations (represented by the pairwise maximum entropy model P_2) and (ii) that large-scale patterns are driven by similarities in people’s weekly routines (represented by the conditionally independent model P_C). Randomly selecting 300 groups of 10 people, Fig. 10(d) shows that the pattern rates predicted by the pairwise maximum entropy model are tightly correlated with the observed pattern rates, avoiding the inaccuracies of the independent and conditionally independent models.

Additionally, calculating the Jensen-Shannon divergences $D_{JS}(Q||P)$ from each model Q to the observed data P , we find that one would typically need over 5 times more samples to distinguish the pairwise model than the independent model [Fig. 10(e)], reflecting roughly the same performance as in the network of email correspondence. Interestingly, in contrast to email activity, the conditionally independent model provides nearly no improvement over the independent model in the data set of private messages. Additionally, Fig. 10(f) shows that the pairwise maximum entropy model captures $I_2/I \approx 87\%$ of the correlation in the data, nearly identical to its performance on the network of email correspondence, while the conditionally independent model accounts for a strikingly small fraction of the correlation structure ($I_C/I \approx 5\%$). This difference in the performance of P_C between the private message and email data sets suggests that the conversational nature of private messages makes them less likely than email traffic to depend on people’s routines. By contrast, the maximum entropy model accurately describes the activity in both populations, further validating the conclusion that patterns of collective behavior can be understood as emerging from simple pairwise correlations.

2. Physical contacts

Thus far, we have only studied human actions mediated by online communication. Here, we instead consider a data set of face-to-face interactions between 50 attendees at the

ACM Hypertext 2009 conference, which spanned 3 days [50]. Discretizing the population activity into bins of width $\Delta t = 20$ seconds, we arrive at a set of about 10^4 binary activity vectors. As in both the networks of email and private message correspondence, we observe that the number of human contacts within a given 20-second window roughly obeys an exponential distribution, while the independent model instead predicts a Poisson distribution that severely underpredicts the likelihood of surges in human activity [Fig. 11(a)]. To study the pairwise maximum entropy model, we generate 200 random groups of 10 individuals. Figure 11(b) shows that the rates of activity patterns predicted by the pairwise model are tightly correlated with the rates at which they were observed at the conference, providing consistently more accurate predictions than both the independent and conditionally independent models.

Quantitatively, one would require 3 to 4 times as many samples to distinguish the independent model from the

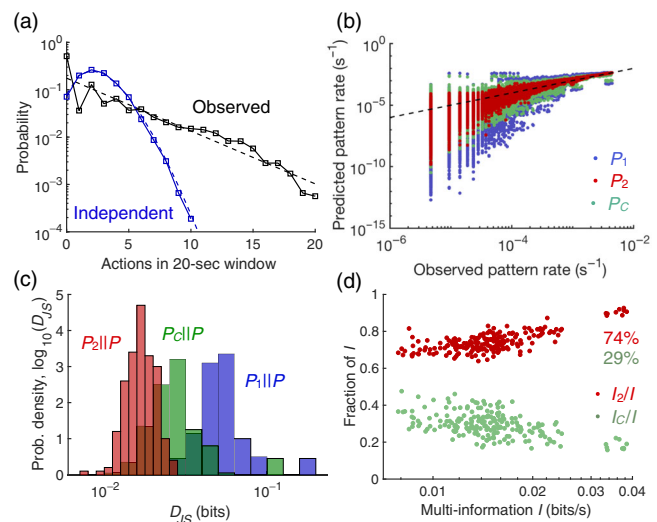


FIG. 11. Performance of the pairwise model in a data set of face-to-face contacts between individuals. (a) Distribution of the number of contacts in a given 20-second window observed in the data set (black) and after shuffling individuals’ activities to eliminate correlations (blue). Dashed lines indicate an exponential fit to the observed data (black) and a Poisson fit to the shuffled data (blue). (b) The rate of each observed activity pattern across 200 randomly selected groups of 10 individuals plotted against the approximate rates under the independent model P_1 (blue), the pairwise maximum entropy model P_2 (red), and the conditionally independent model P_C (green). The dashed line indicates equality. (c) Jensen-Shannon divergences between the true distribution P and the independent P_1 (blue), maximum entropy P_2 (red), and conditionally independent P_C (green) models. The histograms reflect estimates from the 200 10-person groups. (d) Fraction of the network correlation (i.e., multi-information I) captured by the pairwise (red) and conditionally independent (green) models, plotted against the full multi-information; I is divided by $\Delta t = 20$ seconds to remove the dependence on window size.

observed data than the maximum entropy model, and the maximum entropy model achieves a lower Jensen-Shannon divergence from the observed data than from the conditionally independent model across all 200 groups of attendees [Fig. 11(c)]. Additionally, Fig. 11(d) shows that the pairwise model captures $I_2/I \approx 74\%$ of the correlation in the face-to-face contacts. While this is slightly lower than that observed for emails and private messages, we remark that the conditionally independent model only accounts for $I_C/I \approx 29\%$ of the correlation in the data. Interestingly, despite physical interactions representing a quite different mode of human activity from online communication, we still find that patterns of population behavior are well described as arising from pairwise correlations.

3. Music streams

To this point, our analyses have considered modes of human activity that are themselves types of interactions between individuals. It is natural to suspect, therefore, that these activities might be particularly conducive to being described by a pairwise model. To test the ability of the pairwise maximum entropy model to describe other modes of human activity, here we consider a data set of 610 individuals streaming music on the website `last.fm` over the span of one year [51]. Discretizing the streaming activity into bins of width $\Delta t = 150$ seconds (roughly corresponding to the length of an average song), we arrive at a set of about 2×10^5 activity vectors. Considering the number of music streams in a given 150-second window, we notice that the observed distribution is notably not described by an exponential distribution [Fig. 12(a)], which is attributable to the fact that the streaming data are much less sparse than any of the three activities studied previously. Nevertheless, we still find that the observed distribution is heavy tailed relative to the independent Poisson distribution and is characterized by surges of activity where upwards of 50 users are streaming music at a given time.

Randomly selecting 200 groups of 10 users, we show in Fig. 12(b) that the pairwise maximum entropy model provides a much tighter fit of the observed activity pattern rates than either the independent or conditionally independent models. Moreover, by studying the Jensen-Shannon divergences between the different models and the observed distribution of activity patterns, we find that we would need over 6 times as many data samples to distinguish P_2 from P than to distinguish P_1 from P , and over 4 times more samples to distinguish P_C [Fig. 12(c)]. These results are further supported by Fig. 12(d), which shows that the pairwise model captures $I_2/I \approx 74\%$ of the correlation in groups of 10 users, nearly identical to the case of face-to-face contacts. Meanwhile, the daily and weekly rhythms only account for $I_C/I_N \approx 35\%$ of the correlation in the data.

All together, our analysis of private messages, face-to-face contacts, and online music streams serve to strengthen the conclusions made in the main text, namely, that

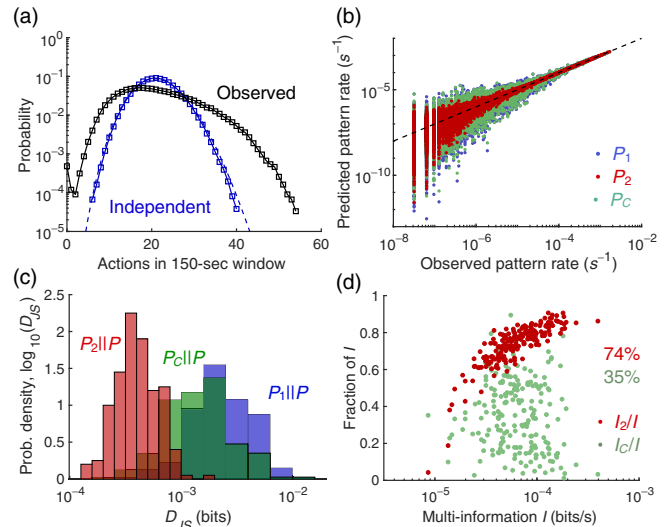


FIG. 12. Performance of the maximum entropy model in a data set of music streams. (a) Distribution of the number of streams in a given 150-second window in the data set (black) and after shuffling individuals' activities to eliminate correlations (blue). The dashed line indicates a Poisson fit to the shuffled data (blue). (b) The rate of each observed activity pattern across 200 randomly selected groups of 10 individuals plotted against the approximate rates under the independent model P_1 (blue), the pairwise maximum entropy model P_2 (red), and the conditionally independent model P_C (green). The dashed line indicates equality. (c) Jensen-Shannon divergences between the true distribution P and the independent P_1 (blue), maximum entropy P_2 (red), and conditionally independent P_C (green) models. The histograms reflect estimates from the 200 10-person groups. (d) Fraction of the network correlation (i.e., multi-information I) captured by the pairwise (red) and conditionally independent (green) models, plotted against the full multi-information. Here, I is divided by $\Delta t = 150$ seconds to remove the dependence on window size.

pairwise correlations can build upon one another to generate predictable patterns of population-wide activity.

APPENDIX D: LEARNING A PAIRWISE MAXIMUM ENTROPY MODEL: THE INVERSE ISING PROBLEM

Here we present the theory and methodology behind learning a pairwise maximum entropy model of collective human activity. Specifically, we describe how to calculate the Ising parameters h_i and J_{ij} from a data set of collective activity patterns. This inference task has a rich history in machine learning under the title Boltzmann machine learning [43] and is commonly referred to in physics as the inverse Ising problem [52].

1. Exact models for small populations

Given the observed distribution P of activity patterns, there is a unique pairwise model P_2 that is consistent with the observed activity rates $\langle \sigma_i \rangle$ and pairwise correlations

$\langle \sigma_i \sigma_j \rangle$, where $\langle \cdot \rangle$ represents an average over P . To learn this pairwise model, one typically begins with an initial pairwise distribution Q with parameters \tilde{h}_i and \tilde{J}_{ij} and then performs gradient descent in the model parameters, with gradients defined by

$$\Delta \tilde{h}_i \propto \langle \sigma_i \rangle - \langle \sigma_i \rangle_Q, \quad (\text{D1})$$

$$\Delta \tilde{J}_{ij} \propto \langle \sigma_i \sigma_j \rangle - \langle \sigma_i \sigma_j \rangle_Q, \quad (\text{D2})$$

where $\langle \cdot \rangle_Q$ represents an average over Q . For groups of size $N = 10$, these gradient calculations are tractable and standard gradient descent converges to the correct pairwise maximum entropy model P_2 .

2. Approximate models for large populations

The primary difficulty in learning a maximum entropy model for a large population, such as the group of 100 email users, lies in calculating the one- and two-point correlations under Q at each gradient step in Eqs. (D1) and (D2). For large populations, exact calculations using the Boltzmann distribution are infeasible, and one must resort to approximate methods. The standard strategy is to simulate the system using Monte Carlo techniques [31,53,54]. Naively, one would run a new Monte Carlo simulation to estimate the gradients at each step of the learning algorithm. However,

this straightforward approach is extremely inefficient. Instead, one can adjust the estimates of the one- and two-point correlations at each gradient step using importance sampling [55] or histogram Monte Carlo [42]. In addition to limiting the number of Monte Carlo simulations, because each sample σ of Q is dominated by inactive individuals, one can leverage sparse matrix operations to significantly speed up the simulations themselves.

We terminate the learning algorithm when the model correlations, $\langle \sigma_i \rangle_Q$ and $\langle \sigma_i \sigma_j \rangle_Q$, are sufficiently close to the observed correlations. The relevant scale for errors in the observed correlations is defined by the standard deviations $\Delta \langle \sigma_i \rangle$ and $\Delta \langle \sigma_i \sigma_j \rangle$, which are estimated by bootstrap sampling from the original data set. Thus, the learning algorithm is terminated when

$$|\langle \sigma_i \rangle - \langle \sigma_i \rangle_Q| < \Delta \langle \sigma_i \rangle \approx 2.2 \times 10^{-4} \quad (\text{D3})$$

and

$$|\langle \sigma_i \sigma_j \rangle - \langle \sigma_i \sigma_j \rangle_Q| < \Delta \langle \sigma_i \sigma_j \rangle \approx 1.7 \times 10^{-4}. \quad (\text{D4})$$

We confirm that the individual email rates and pairwise correlations under the maximum entropy model P_2 match the observed correlations within the experimental errors in the data [Figs. 13(a)–13(c)].

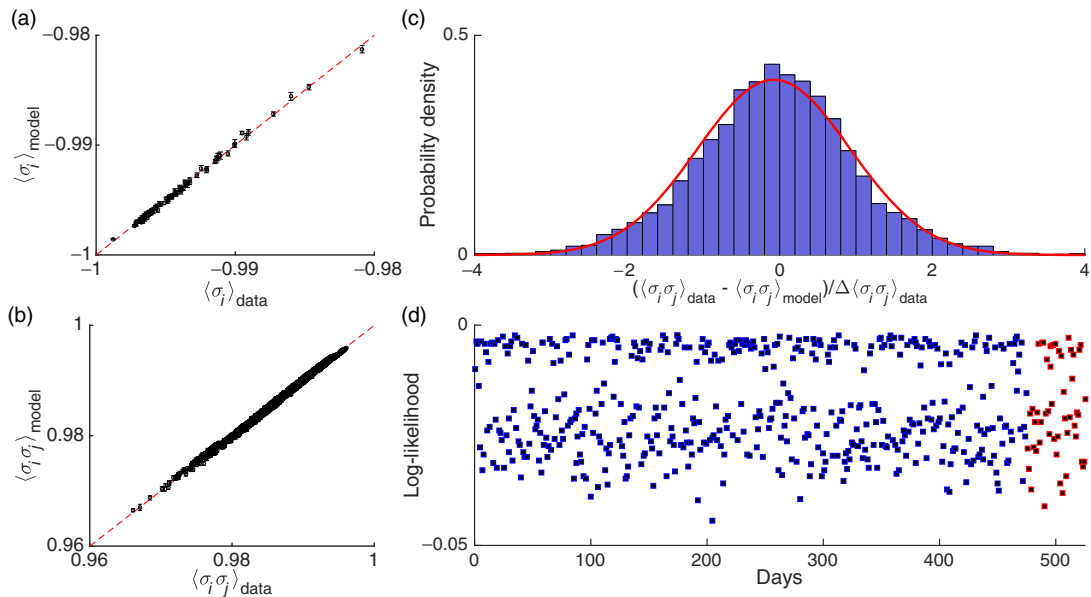


FIG. 13. Learning a pairwise maximum entropy model for a 100-person population. (a) Reconstructed activity rates for all 100 individuals under the maximum entropy model, plotted against their true activity rates. The dashed line indicates equality. (b) Reconstructed pairwise correlations under the maximum entropy model versus the observed correlations. (c) Distribution of the differences between the true and model pairwise correlations, normalized by the error in the data $\Delta \langle \sigma_i \sigma_j \rangle$. For reference, the red line is a Gaussian distribution with unit variance. The empirically measured distribution has nearly Gaussian shape with a standard deviation of approximately 1.05, demonstrating that the learning algorithm reconstructs the pairwise correlations within experimental precision. (d) The per-person average log-likelihood of the data $\langle \log P_2(\sigma) \rangle / N$, where the average is taken over all patterns within a given day, computed for the training days (blue) and test days (red). The data have been sorted so that the test days follow the training days, but the true choice of test days was random.

For a population of 100 individuals, defining a pairwise maximum entropy model requires learning $N(N+1)/2 = 5050$ different parameters. Given such a large number, it is possible that the model is being finely tuned to match statistical errors in the data. To test for overfitting, we randomly select 476 of the 526 days to learn the model, and then we test the accuracy of the model on the remaining 50 days. We confirm that the pairwise model assigns the same amount of probability to the test data as to the training data, within errors, demonstrating that the learned model generalizes to describe data outside of the training set [Fig. 13(d)]. We conclude that the learned pairwise model (i) fits the activity data within experimental precision and (ii) does not overfit statistical noise in the data. For access to the calculated external fields h_i and pairwise interactions J_{ij} , please contact the corresponding author.

APPENDIX E: THE CONDITIONALLY INDEPENDENT MODEL

To test the prediction that collective behavior is driven by similarities in people's daily and weekly routines, we study the conditionally independent model P_C . Letting $p_i^t(\sigma_i)$ denote the probability of person i performing an action within a window of width Δt at time t during the week, the conditionally independent model is defined by

$$P_C(\sigma) = \frac{\Delta t}{\omega} \sum_t \prod_i p_i^t(\sigma_i), \quad (\text{E1})$$

where ω denotes the length of a day or week. Under this conditionally independent model, correlations between individuals are driven by fluctuations in their inherent activity rates.

APPENDIX F: ESTIMATING ENTROPY FROM FINITE DATA

To calculate the multi-information $I = S_1 - S$ of the network activity, we must compute the entropies of the independent model S_1 and the observed data S . While calculating S_1 is straightforward, we must estimate the true entropy S from a finite number of samples, possibly leading to finite-size errors. Suppose that the data set consists of the patterns $\{\sigma^\alpha\}$ with corresponding probabilities $\{p^\alpha\}$. One could naively estimate the entropy using the standard formula

$$\tilde{S} = -\sum_\alpha p^\alpha \log p^\alpha. \quad (\text{F1})$$

However, since some of the patterns are likely missing and the probabilities p^α are not exact, this estimate should fundamentally be viewed as an approximation to S that improves as the number of samples increases. To correct for the sample size dependence of \tilde{S} , we subsample the data

and fit the resulting estimates using a form proposed by Strong *et al.* [56],

$$\tilde{S}(\text{size}) = S + \frac{a}{\text{size}} + \frac{b}{\text{size}^2}, \quad (\text{F2})$$

where a and b are finite-size corrections. Using this fit, we can extract an accurate estimate of the true entropy S . We remark that for large data sets such as those considered here, and for relatively small networks like the 10-person groups studied in the main text, finite-size errors are small.

APPENDIX G: EXTENDED DISCUSSION

Our investigation of collective human behavior yields three distinct conclusions:

- (1) Large-scale behavior, characterized by surges in collective activity, cannot be understood using models that assume humans behave independently.
- (2) While collective behavior is far from independent, the minimal extension of the independent model consistent with the observed pairwise correlations captures most of the correlation in all populations considered, accurately predicting surges of collective activity.
- (3) In the network of email correspondence, the learned pairwise interactions are closely related to the underlying topology of interhuman communication, imbuing the maximum entropy model with real-world interpretability.

Here, we discuss the implications and limitations of these results while keeping in mind that modern life involves a diverse range of activities, some of which may require a fundamentally different approach. Throughout, we emphasize important opportunities for future research.

1. Internal correlations versus external influences

In the study of human dynamics, as in the study of physical and biological systems, any macroscopic behavior that evades explanation by a model of independent elements fundamentally derives from two possible sources of correlation: (i) interactions between elements or groups of elements and (ii) external influences on the system. In all human activities considered here, we witness surges of collective activity that cannot be explained under assumptions of human independence. Instead, we find that the populations are described quantitatively by models that include the simplest possible correlations—those between pairs of individuals. However, given that large-scale patterns could derive from higher-order correlations or from shared external inputs to the population, and given the myriad experiences that shape human actions, it would be naive to universally conclude that all collective human activity emerges from pairwise correlations. Instead, we hypothesize that particular activities fall along a spectrum,

with internal correlations and external influences each playing roles of variable importance.

We remark that we have already witnessed evidence for such a spectrum in the different human activities studied above. For example, while patterns of email communication were reasonably well described by taking into account people's weekly rhythms, capturing about 67% of the correlation structure in 10-person groups, private message correspondence had a markedly weak dependence on people's schedules, with daily routines accounting for only about 5% of the correlation in 10-person groups. These results agree with intuition, indicating that email activity is moderately tied to people's work and leisure schedules, while daily routines have nearly no predictive power in a network of private messages. Interestingly, correlations in both face-to-face contacts and online music streaming are moderately driven by daily and weekly routines, falling in between email and private message correspondence. With these results in mind, the clearest direction for future investigation is to continue probing different ends of the spectrum by quantifying the relative importance of internal correlations versus external influences in different modes of human behavior.

2. The energy landscape of collective human behavior

Every maximum entropy model Q is defined by a Boltzmann distribution $Q(\boldsymbol{\sigma}) = \exp(-E(\boldsymbol{\sigma}))/Z$, where $E(\boldsymbol{\sigma})$ is the energy function, or Hamiltonian, that describes the system, and Z is the normalization constant. In the case of the pairwise maximum entropy model, the relevant energy function is that of the Ising model, $E(\boldsymbol{\sigma}) = -\frac{1}{2} \sum_{i \neq j} J_{ij} \sigma_i \sigma_j - \sum_i h_i \sigma_i$. In statistical mechanics, there is a wealth of literature exploring the diversity of large-scale behaviors that can emerge from systems with different energy landscapes [33,57]. Thus, future research should leverage this connection to answer a number of important questions: What can the energy landscape of a given population tell us about its functional properties? Does collective human behavior favor dramatic shifts in activity, or are social populations organized to incentivize local fluctuations, guarding against the effects of large external shocks?

3. Beyond equal-time correlations

Throughout our analysis, we have focused on modeling equal-time correlations, which quantify the tendencies of individuals to engage in synchronous actions. In doing so, we have implicitly assumed that each observed activity pattern $\boldsymbol{\sigma}$ is drawn independently from an underlying stationary distribution $P(\boldsymbol{\sigma})$, leaving models of the population's activity without notions of time or causality. While studying equal-time correlations has allowed us to reach a number of important conclusions, the idea that patterns of human activity are sampled from a stationary

distribution is not consistent with the common intuition that conscious human actions are often responses to prior social and environmental influences. For example, the fact that individuals perform bursts of actions in quick succession is thought to be the result of a decision-based queuing process [1], and it is known that the temporal scales of human activity can change over time [58–60].

In the context of human communication, a significant fraction of emails and private messages are direct responses to previous correspondence. Therefore, it would be interesting to study the correlations between people's activities with a time delay τ in between, where τ represents the characteristic response time of communication in the population. Such spatiotemporal correlations have recently received a large amount of interest in neuroscience and biology, where it has been found that the spatiotemporal patterns of spiking neurons in the brain and flocks of birds in flight are only partially captured by stationary maximum entropy models [32,61,62]. Similarly, studying the spatiotemporal patterns that define collective human activity has significant implications for understanding the causal flow of influences and information between individuals in a population [60]. Furthermore, developing accurate dynamical models of large-scale behavior has important ramifications for predicting the effects of interventions and time-varying perturbations in networks of interacting humans [5,7,8,24,25,63,64].

-
- [1] A.-L. Barabási *The Origin of Bursts and Heavy Tails in Human Dynamics*, *Nature (London)* **435**, 207 (2005).
 - [2] A. Vázquez, J. G. Oliveira, Z. Dezső, K.-I. Goh, I. Kondor, and A.-L. Barabási, *Modeling Bursts and Heavy Tails in Human Dynamics*, *Phys. Rev. E* **73**, 036127 (2006).
 - [3] D. Rybski, S. V. Buldyrev, S. Havlin, F. Liljeros, and H. A. Makse, *Scaling Laws of Human Interaction Activity*, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 12640 (2009).
 - [4] J.-P. Eckmann, E. Moses, and D. Sergi, *Entropy of Dialogues Creates Coherent Structures in E-mail Traffic*, *Proc. Natl. Acad. Sci. U.S.A.* **101**, 14333 (2004).
 - [5] C. Peng, X. Jin, K.-C. Wong, M. Shi, and P. Liò, *Collective Human Mobility Pattern from Taxi Trips in Urban Area*, *PLoS One* **7**, e34487 (2012).
 - [6] J. P. Bagrow, D. Wang, and A.-L. Barabasi, *Collective Response of Human Populations to Large-Scale Emergencies*, *PLoS One* **6**, e17680 (2011).
 - [7] R. Crane and D. Sornette, *Robust Dynamic Classes Revealed by Measuring the Response Function of a Social System*, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 15649 (2008).
 - [8] J. Candia, M. C. González, P. Wang, T. Schoenharl, G. Madey, and A.-L. Barabási *Uncovering Individual and Collective Human Dynamics from Mobile Phone Records*, *J. Phys. A* **41**, 224015 (2008).
 - [9] C. Castellano, S. Fortunato, and V. Loreto, *Statistical Physics of Social Dynamics*, *Rev. Mod. Phys.* **81**, 591 (2009).

- [10] M. Perc, J. J. Jordan, D. G. Rand, Z. Wang, S. Boccaletti, and A. Szolnoki, *Statistical Physics of Human Cooperation*, *Phys. Rep.* **687**, 1 (2017).
- [11] S. Galam, Y. Gefen, and Y. Shapir, *Sociophysics: A New Approach of Sociological Collective Behaviour. I. Mean-Behaviour Description of a Strike*, *J. Math. Sociol.* **9**, 1 (1982).
- [12] S. Galam and S. Moscovici, *Towards a Theory of Collective Phenomena: Consensus and Attitude Changes in Groups*, *Eur. J. Soc. Psychol.* **21**, 49 (1991).
- [13] Z. Wang, C. T. Bauch, S. Bhattacharyya, A. d'Onofrio, P. Manfredi, M. Perc, N. Perra, M. Salathé, and D. Zhao, *Statistical Physics of Vaccination*, *Phys. Rep.* **664**, 1 (2016).
- [14] A.-L. Barabási, H. Jeong, Z. Néda, E. Ravasz, A. Schubert, and T. Vicsek, *Evolution of the Social Network of Scientific Collaborations*, *Physica A* **311**, 590 (2002).
- [15] M. R. D'Orsogna and M. Perc, *Statistical Physics of Crime: A Review*, *Phys. Life Rev.* **12**, 1 (2015).
- [16] D. Helbing, D. Brockmann, T. Chadeaux, K. Donnay, U. Blanke, O. Woolley-Meza, M. Moussaid, A. Johansson, J. Krause, S. Schutte *et al.*, *Saving Human Lives: What Complexity Science and Information Systems Can Contribute*, *J. Stat. Phys.* **158**, 735 (2015).
- [17] F. A. Haight, *Handbook of the Poisson Distribution* (Wiley, New York, 1967).
- [18] G. Rasch, *The Poisson Process as a Model for a Diversity of Behavioral Phenomena*, in *International Congress of Psychology*, Vol. 2 (American Psychological Association, Washington DC, 1963), p. 2.
- [19] T. Karagiannis, M. Molle, and M. Faloutsos, *Long-Range Dependence Ten Years of Internet Traffic Modeling*, *IEEE Internet Comput.* **8**, 57 (2004).
- [20] D. L. Gerlough and A. Schuhl, *Use of Poisson Distribution in Highway Traffic* (Eno Foundation for Highway Traffic Control, Saugatuck, CT, 1955).
- [21] S. Galam, *Rational Group Decision Making: A Random Field Ising Model at $T = 0$* , *Physica A* **238**, 66 (1997).
- [22] S. Galam, *Sociophysics: A Review of Galam Models*, *Int. J. Mod. Phys. C* **19**, 409 (2008).
- [23] C. W. Lynn and D. D. Lee, *Statistical Mechanics of Influence Maximization with Thermal Noise*, *Europhys. Lett.* **117**, 66001 (2017).
- [24] J.-P. Onnela, J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A.-L. Barabási, *Structure and Tie Strengths in Mobile Communication Networks*, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 7332 (2007).
- [25] R. Pastor-Satorras and J. Vespignani, *Epidemic Spreading in Scale-free Networks*, *Phys. Rev. Lett.* **86**, 3200 (2001).
- [26] A. Paranjape, A. R. Benson, and J. Leskovec, *Motifs in Temporal Networks*, in *Proc. ACM WSDM* (ACM, Cambridge, England, 2017), pp. 601–610.
- [27] P. Panzarasa, T. Opsahl, and K. M. Carley, *Patterns and Dynamics of Users' Behavior and Interaction: Network Analysis of an Online Community*, *J. Assoc. Inf. Sci. Technol.* **60**, 911 (2009).
- [28] D. Sornette, F. Deschâtres, T. Gilbert, and Y. Ageon, *Endogenous Versus Exogenous Shocks in Complex Networks: An Empirical Test Using Book Sale Rankings*, *Phys. Rev. Lett.* **93**, 228701 (2004).
- [29] F. Deschâtres and D. Sornette, *Dynamics of Book Sales: Endogenous Versus Exogenous Shocks in Complex Networks*, *Phys. Rev. E* **72**, 016112 (2005).
- [30] E. Schneidman, M. J. Berry, R. Segev II, and W. Bialek, *Weak Pairwise Correlations Imply Strongly Correlated Network States in a Neural Population*, *Nature (London)* **440**, 1007 (2006).
- [31] E. Ganmor, R. Segev, and E. Schneidman *Sparse Low-Order Interaction Network Underlies a Highly Correlated and Learnable Neural Population Code*, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 9679 (2011).
- [32] O. Marre, S. El Boustani, Y. Frégnac, and A. Destexhe, *Prediction of Spatiotemporal Patterns of Neural Activity from Pairwise Correlations*, *Phys. Rev. Lett.* **102**, 138101 (2009).
- [33] E. T. Jaynes, *Information Theory and Statistical Mechanics*, *Phys. Rev.* **106**, 620 (1957).
- [34] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (John Wiley & Sons, New York, 2012).
- [35] W. Bialek, A. Cavagna, I. Giardina, T. Mora, E. Silvestri, M. Viale, and A. M. Walczak, *Statistical Mechanics for Natural Flocks of Birds*, *Proc. Natl. Acad. Sci. U.S.A.* **109**, 4786 (2012).
- [36] M. Weigt, R. A. White, H. Szurmant, J. A. Hoch, and T. Hwa, *Identification of Direct Residue Contacts in Protein-Protein Interaction by Message Passing*, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 67 (2009).
- [37] T. R. Lezon, J. R. Banavar, M. Cieplak, A. Maritan, and N. V. Fedoroff, *Using the Principle of Entropy Maximization to Infer Genetic Interaction Networks from Gene Expression Patterns*, *Proc. Natl. Acad. Sci. U.S.A.* **103**, 19033 (2006).
- [38] C. W. Lynn and D. D. Lee, *Maximizing Influence in an Ising Network: A Mean-Field Optimal Solution*, in *Advances in Neural Information Processing Systems* (Curran Association, Red Hook, NY, 2016), pp. 2495–2503.
- [39] J. Lin, *Divergence Measures Based on the Shannon Entropy*, *IEEE Trans. Inf. Theory* **37**, 145 (1991).
- [40] R. D. Malmgren, D. B. Stouffer, A. E. Motter, and L. A. N. Amaral, *A Poissonian explanation for heavy tails in e-mail communication*, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 18153 (2008).
- [41] H. Barlow, *Conditions for Versatile Learning, Helmholtz's Unconscious Inference, and the Task of Perception*, *Vision Res.* **30**, 1561 (1990).
- [42] A. M. Ferrenberg and R. H. Swendsen, *New Monte Carlo Technique for Studying Phase Transitions*, *Phys. Rev. Lett.* **61**, 2635 (1988).
- [43] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, *A Learning Algorithm for Boltzmann Machines*, *Cogn. Sci.* **9**, 147 (1985).
- [44] M. Mézard, G. Parisi, and M. Virasoro, *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications* (World Scientific, Singapore, 1987), Vol. 9.
- [45] S. Galam, *Fragmentation Versus Stability in Bimodal Coalitions*, *Physica A* **230**, 174 (1996).
- [46] A. Ashourvan, S. Gu, M. G. Mattar, J. M. Vettel, and D. S. Bassett, *The Energy Landscape Underpinning Module Dynamics in the Human Brain Connectome*, *NeuroImage* **157** 364 (2017).

- [47] D. Kempe, J. Kleinberg, and É. Tardos, *Maximizing the Spread of Influence through a Social Network*, in *SIGKDD* (ACM, Washington DC, 2003), pp. 137–146.
- [48] F. Morone and H. A. Makse, *Influence Maximization in Complex Networks through Optimal Percolation*, *Nature (London)* **524**, 65 (2015).
- [49] C. W. Lynn and D. D. Lee, *Maximizing Activity in Ising Networks via the TAP Approximation*, in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, (AAAI Press, San Jose, CA, 2018), pp. 679–686.
- [50] L. Isella, J. Stehlé, A. Barrat, C. Cattuto, J.-F. Pinton, and W. Van den Broeck, *What's in a Crowd? Analysis of Face-to-Face Behavioral Networks*, *J. Theor. Biol.* **271**, 166 (2011).
- [51] O. Celma, *Music Recommendation*, in *Music Recommendation and Discovery* (Springer, New York, 2010), pp. 43–85.
- [52] E. Aurell and M. Ekeberg, *Inverse Ising Inference Using All the Data*, *Phys. Rev. Lett.* **108**, 090201 (2012).
- [53] W. R. Gilks, S. Richardson, and D. Spiegelhalter, *Markov Chain Monte Carlo in Practice* (CRC Press, New York, NY, 1995).
- [54] G. Tkačik, O. Marre, D. Amodei, E. Schneidman, W. Bialek, and M. J. Berry II, *Searching for Collective Behavior in a Large Network of Sensory Neurons*, *PLoS Comput. Biol.* **10**, e1003408 (2014).
- [55] M. I. Jordan, *Learning in Graphical Models* (Springer Science & Business Media, New York, 1998), Vol. 89.
- [56] S. P. Strong, R. Koberle, R. R. de Ruyter van Steveninck, and W. Bialek, *Entropy and Information in Neural Spike Trains*, *Phys. Rev. Lett.* **80**, 197 (1998).
- [57] D. Chandler, *Introduction to Modern Statistical Mechanics* (Oxford University, New York, 1987).
- [58] G. Miritello, R. Lara, M. Cebrian, and E. Moro, *Limited Communication Capacity Unveils Strategies for Human Interaction*, *Sci. Rep.* **3**, 1950 (2013).
- [59] J. Borge-Holthoefer, N. Perra, B. Gonçalves, S. González-Bailón, A. Arenas, Y. Moreno, and A. Vespignani, *The Dynamics of Information-Driven Coordination Phenomena: A Transfer Entropy Analysis*, *Sci. Adv.* **2**, e1501158 (2016).
- [60] H. Navarro, G. Miritello, A. Canales, and E. Moro, *Temporal Patterns Behind the Strength of Persistent Ties*, *EPJ Data Science* **6**, 31 (2017).
- [61] A. Tang, D. Jackson, J. Hobbs, W. Chen, J. L. Smith, H. Patel, A. Prieto, D. Petrusca, M. I. Grivich, A. Sher *et al.*, *A Maximum Entropy Model Applied to Spatial and Temporal Correlations from Cortical Networks in vitro*, *J. Neurosci.* **28**, 505 (2008).
- [62] A. Cavagna, I. Giardina, F. Ginelli, T. Mora, D. Piovani, R. Tavarone, and A. M. Walczak, *Dynamical Maximum Entropy Approach to Flocking*, *Phys. Rev. E* **89**, 042707 (2014).
- [63] M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi, *Understanding Individual Human Mobility Patterns*, *Nature (London)* **453**, 779 (2008).
- [64] G. Caldarelli, M. Marsili, and Y.-C. Zhang, *A Prototype Model of Stock Exchange*, *Europhys. Lett.* **40**, 479 (1997).

Correction: A proof change request was not fully implemented in the fourth sentence of the last paragraph of the Introduction and has now been set right.