INFORMATION SCIENCE

10.1126/sciadv.1501158

The dynamics of information-driven coordination phenomena: A transfer entropy analysis

Javier Borge-Holthoefer,[1*†] Nicola Perra,[2*] Bruno Gonçalves,[3‡] Sandra González-Bailón,[4] Alex Arenas,[5*] Yamir Moreno,[6,7,8*] Alessandro Vespignani[2,8,9*]

Data from social media provide unprecedented opportunities to investigate the processes that govern the dynamics of collective social phenomena. We consider an information theoretical approach to define and measure the temporal and structural signatures typical of collective social events as they arise and gain prominence. We use the symbolic transfer entropy analysis of microblogging time series to extract directed networks of influence among geolocalized subunits in social systems. This methodology captures the emergence of system-level dynamics close to the onset of socially relevant collective phenomena. The framework is validated against a detailed empirical analysis of five case studies. In particular, we identify a change in the characteristic time scale of the information transfer that flags the onset of information-driven collective phenomena. Furthermore, our approach identifies an order-disorder transition in the directed network of influence between social subunits. In the absence of clear exogenous driving, social collective phenomena can be represented as endogenously driven structural transitions of the information transfer network. This study provides results that can help define models and predictive algorithms for the analysis of societal events based on open source data.

INTRODUCTION

A vivid scientific and popular media debate has recently centered on the role that social networking tools play in coordinating collective phenomena; examples include street protests, civil unrests, consensus formation, and the emergence of electoral preferences. A flurry of studies have analyzed the correlation of search engine queries, microblogging posts, and other open data sources with the incidence of infectious disease (*1*–*4*), box office returns (*5*), stock market behavior (*6*, *7*), election outcomes (*8*, *9*), popular votes results (*10*), crowd sizes (*11*), and social unrest (*12*, *13*). However, many other studies have also pointed out the challenges that big data present and the likely methodological pitfalls that might result from their analysis (*14*–*19*). Previous work suggests that more research is needed to develop methods for exploiting the value of social media data while overcoming their limitations.

Here, we use microblogging data to extract networks of causal influence among different geographical subunits before, during, and after collective social phenomena. To ground our work on empirical data, we analyze five data sets that track Twitter communication around five well-known social events: the release of a Hollywood blockbuster movie, two massive political protests, the discovery of the Higgs boson, and the acquisition of Motorola by Google. We selected these case studies because they represent different points in a theoretical continuum that separates two types of collective phenomena: those that can be represented as an endogenously driven exchange of information and those that respond more clearly to factors that are exogenous to the system. In our context, these phenomena refer to the dynamics of information exchange through social media. In some cases, discussions evolve organically, building up momentum up to the point where the exchange of information is generalized; however, in some other cases, the discussions emerge suddenly as a reaction to some unexpected external event (*20*). The Motorola-Google case study corresponds to the exogenous type, providing a counterexample and intuitive baseline test for the other cases.

For each case study, we adopt the transfer entropy approach to define an effective social connectivity at the macroscale and study the coordinated activation of localized populations. We address two foundational problems: first, the identification of the characteristic time scale of social events as they develop, gather force, and burst into the public eye. In effect, the determination of the pertinent time scale is one of the fundamental limitations of the analysis of data from social media, namely, to be considered for the posterior processing of information. Symbolic transfer entropy (STE) captures the intrinsic time scale of the information flow and allows a proper diagnosis of the granularity needed to grasp the evolution of social events. Second, we look into the characterization of the structural signature typical of the communication dynamics that underlie social phenomena. We find that the onset of collective phenomena is characterized by the drop of the characteristic time scale; we also show that the emergence of coherent patterns of information flow can be mapped into order-disorder transitions in the underlying connectivity patterns of the transfer entropy network. The methodology we present here can

[1]Qatar Computing Research Institute, Hamad Bin Khalifa University, P.O. Box 5825 Doha, Qatar. [2]Laboratory for the Modeling of Biological and Socio-Technical Systems, Northeastern University, Boston, MA 02115, USA. [3]Aix-Marseille Université, Université de Toulon, CNRS, CPT, UMR 7332, 13288 Marseille, France. [4]Annenberg School for Communication, University of Pennsylvania, Philadelphia, PA 19104, USA. [5]Departament d'Enginyeria Informàtica i Matemàtiques, Universitat Rovira i Virgili, 43007 Tarragona, Spain. [6]Department of Theoretical Physics, Faculty of Sciences, University of Zaragoza, 50009 Zaragoza, Spain. [7]Institute for Biocomputation and Physics of Complex Systems (BIFI), University of Zaragoza, 50018 Zaragoza, Spain. [8]Institute for Scientific Interchange, 10126 Torino, Italy. [9]Institute for Quantitative Social Sciences at Harvard University, Cambridge, MA 02138, USA.
*Corresponding author. E-mail: borge.holthoefer@gmail.com (J.B.-H.); nicolaperra@gmail.com (N.P.); alexandre.arenas@urv.cat (A.A.); yamir.moreno@gmail.com (Y.M.); a.vespignani@neu.edu (A.V.)
†Present address: Internet Interdisciplinary Institute, Universitat Oberta de Catalunya, 08018 Barcelona, Catalonia, Spain.
‡Present address: Center for Data Science, New York University, New York, NY 10003, USA.

therefore be used to gain new insights on the structural and functional relations occurring in large-scale structured populations, eventually leading to the identification of metrics that might be used for the definition of precursors of large-scale social events.

RESULTS

Our data sets consist of time-stamped and geolocalized time series of tweets associated to the following events: the Spanish 15M social unrest in 2011, the Outono Brasileiro ("Brazilian Autumn") in 2013, the discovery of the Higgs boson in 2012, the release of a Hollywood blockbuster in 2012, and the acquisition of Motorola by Google in 2011.

The spatiotemporal annotation of each tweet in the time series allows the construction of spatially localized activity maps that help identify, as time unfolds, the role that different geographical subunits played in the global exchange of information. For each data set, the definition of the corresponding spatial unit is performed according to administrative and geographical boundaries as specified in Materials and Methods (see Fig. 1 as well). Note that the map only shows how the signal increases in all regions and does not provide evidence of any unexpected transition, pointing out that volume alone is not a good indicator of the evolution of the events.

The time-stamped series of tweets that originated from each spatial subunit X (supra-urban aggregates) defines the activity time series X_t of the corresponding subunit in the social system. Time stamps are modified for each data set to account for different time zones (see the Supplementary Materials for details). Activity time series encode the role of each geographical subunit, a sort of "who steers whom," and several techniques can be used to detect directed exchange of information across the social system. Here, we characterize the dominating direction of information flow between spatial subunits using STE (*21*, *22*). This well-established technique has been used to infer directional influence between dynamical systems (*23*–*25*) and to analyze patterns of brain connectivity (*26*).

STE quantifies the directional flow of information between two time series, X and Y, by first categorizing the signals in a small set of symbols or alphabet (see section B.3 in the Supplementary Materials), and then computing the joint and conditional probabilities of the sequence indices from the relative frequency of symbols in each sequence, $\hat{X}$ and $\hat{Y}$, as

$$T_{Y,X} = \Sigma p\left(\hat{x}_{i+\delta}, \hat{x}_i, \hat{y}_i\right) \log_2 \left(\frac{p(\hat{x}_{i+\delta} | \hat{x}_i, \hat{y}_i)}{p(\hat{x}_{i+\delta} | \hat{x}_i)} \right) \quad (1)$$

where the sum runs over each symbol in the sequence and $\delta = 1$. The transfer entropy refers to the deviations of the cross-Markovian property of the series (independence between them), measured as the Kullback-Leibler divergence (*27*) (see the Supplementary Materials for all technical details). An important feature of symbolic approaches is that they discount the relative magnitude of each time series; this is important in our case because different geographical units differ largely in population density or Internet penetration rates. Flattening and discretizing the original signal are a key feature to enhancing the sensitivity of our proposal to any changes (even minor ones) in the interaction dynamics between subsystems (see section F in the Supplementary Materials).

Within this framework, we first analyze the temporal patterns characterizing the flow of information. Admittedly, microblogging data can be sampled at different time scales, Δt. To select the optimal sampling rate, we consider all possible pairs (X,Y) of geographical units and measure the total STE in the system $T = \Sigma_{XY} T_{X,Y}$ as a function of Δt. We consider the system-wide characteristic sampling time scale τ as that which maximizes the total information flow T. This quantity provides an indication of the time scale at which the information is being exchanged in the system, which is not necessarily correlated with volume; see the early stages of (A) to (D) in Fig. 2 where the time scale drops by ~50% in a very scarce volume scenario (see section F in the Supplementary Materials as well). The characteristic time scale τ changes as the phenomena under analysis unfold; that is, it decreases as the system approaches the exponential increase in overall activity that signals the onset of the collective event. As shown in the top panels of Fig. 2, τ is a proxy for the internally generated coordination in the system that culminates simultaneously with the time

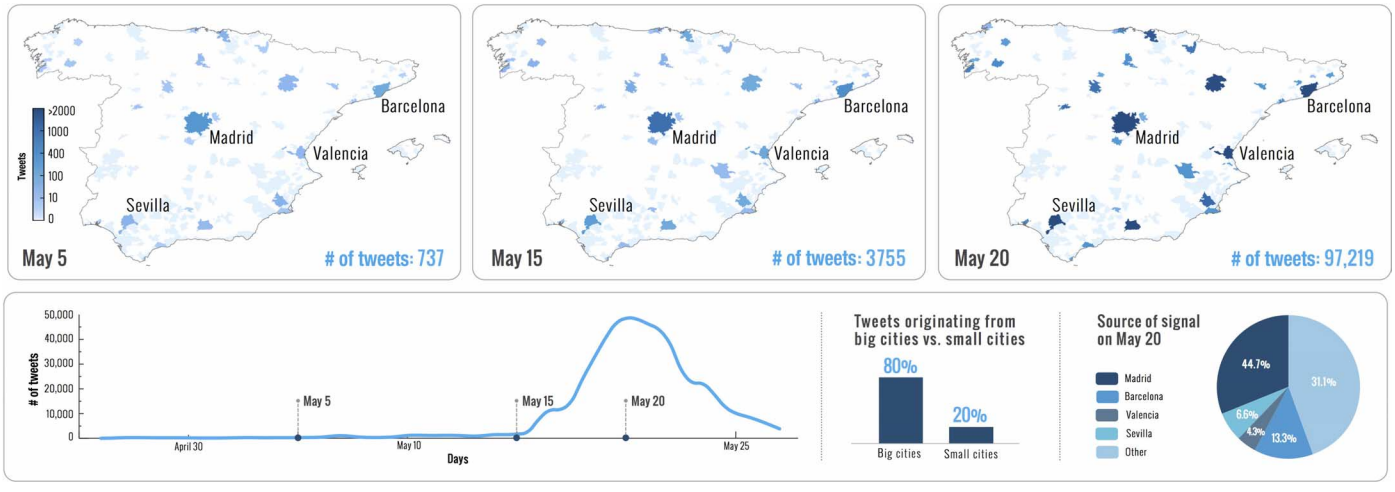

Fig. 1. Spatiotemporal activity as observed from the microblogging platform Twitter. Spain's 15M protest growth in time shows that the protest did not transcend the online sphere until May 15 when the political movement emerged on the streets. Broadcasting traditional media started reporting about it soon after; by that time, demonstrations had been held in the most important cities of the country.

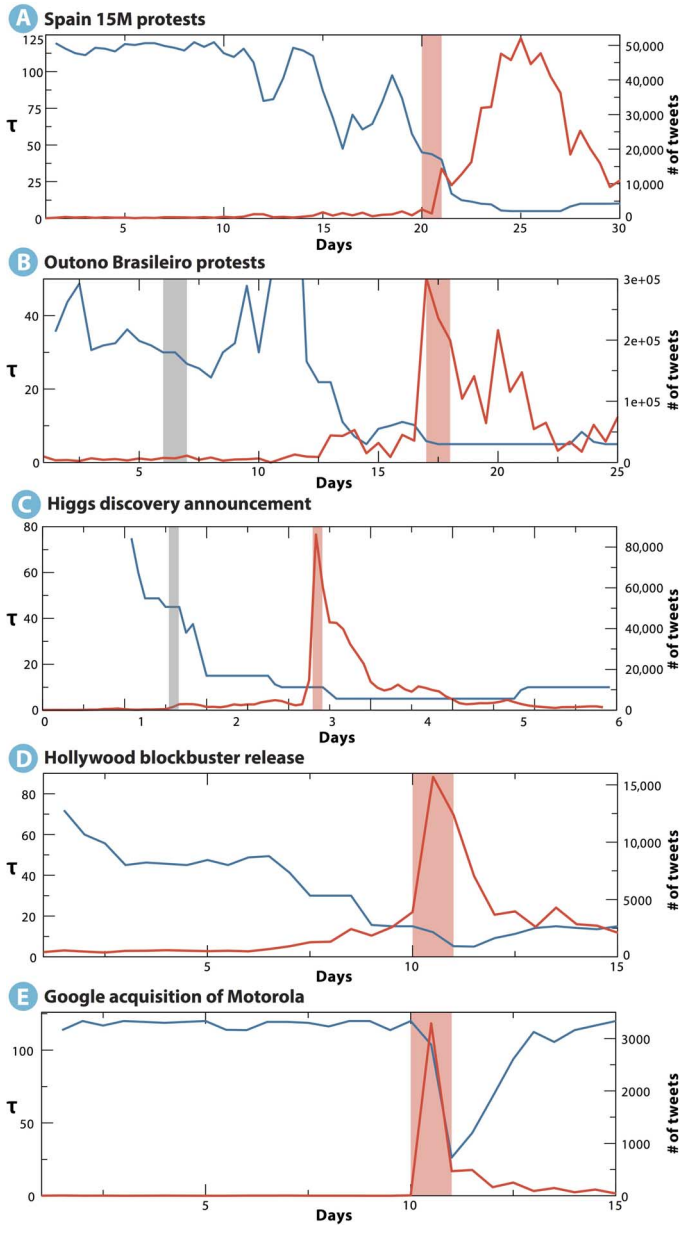

Fig. 2. Characteristic time scale τ. (**A** to **E**) The panels report the variation of the characteristic time scale (blue) that maximizes the STE flow as the social event is approached. Red lines correspond to activity volume (number of tweets). Light red vertical lines correspond to the onset of the main social event. Gray vertical lines (B and C) indicate a smaller precursor event. (A) The 15M event shows a progressive decline of the characteristic time scale well before the actual social event; the same is observed for the Outono Brasileiro in (B) (note a data blackout between days 10 and 11). The patterns for the Higgs boson discovery data set in (C) and the Hollywood blockbuster data (D) also reveal a drop in the characteristic time scale, although this is smoother in the movie case. Overall, in (A) to (D) (endogenous activity), the time scale has already dropped to 50% by the time the absolute volume signals a system-wide event. Finally, the Google-Motorola deal triggers a high volume of microblogging activity without actual change in the time scale of the information flow (E). In this case, the decline is observed in the aftermath of the announcement. As discussed in the main text, this event is the only one that is clearly elicited by an exogenous trigger.

of the occurrence of the social event: the street protest day in the case of political unrest, the movie release date in the case of the Hollywood blockbuster, and the announcement to the press of the Higgs boson discovery. The only clear exception to this behavior is offered by the company acquisition data set: the Google-Motorola announcement is a clear example of collective phenomena driven mostly by an exogenous factor, that is, a media announcement. In this case, the dynamical time scale is constant until the announcement is made public. In the Supplementary Materials, we present the same analysis for the randomized signals, showing that time scale variations are then washed out from the signal, as expected.

The maximized information exchange can be analyzed at the level of geographical subunits by constructing the effective directed network (*28*) of information flow on a daily basis. This network is encoded in the matrix $\{T_{XY}\}$, which contains pairwise information about how each component in the system controls (or is controlled by) the others. The matrix $\{T_{XY}\}$ is asymmetric. The directionality is crucial and denotes that the geographic area x can exert some driving on area y and, at the same time, y might exert some driving on x. For this reason, it is convenient to define the directionality index $T^S_{X,Y} = T_{Y,X} - T_{X,Y}$, which measures the balance of information flow in both directions. This index quantifies the dominant direction of information flow and is expected to have positive values for unidirectional couplings with x as the driver and negative values if y is driving x. For symmetric bidirectional couplings, we expect $T^S_{X,Y}$ to be null.

We show in Fig. 3 the temporal evolution of the maximized $\Sigma_Y T^S_{X,Y}$ that provides the information flow balance of each specific geographical area. The results show that in the 15M grassroots protests, a limited number of urban areas initially drive the onset of the social phenomenon. These units correspond mostly to major cities; however, the analysis also uncovers hidden drivers, such as Orotava, a less known urban area. Only after the first demonstration day on May 15 does the driving role become much more homogeneously distributed. In the Brazilian case, a set of clear drivers is present only during the onset phase preceding a demonstration on June 6 and becomes fuzzier up until the major demonstration (June 17) and totally blurred afterward. We find a similar behavior in the Higgs boson cases (with rumors around the discovery on July 2 and the final announcement on July 4) (*29*). The blockbuster case is driven by a steady excitement of the public before the movie release. Again, as expected, we observe completely different patterns in the case of the Google data set.

In general, the evolving effective networks reveal a transition from hierarchical causal relationships to symmetric (if rather fluctuating) networks where information is flowing symmetrically among all subunits. If information flows mainly in one direction (that is, if the subsystems are arranged in a highly hierarchical structure), a subunit dominates another, with no or little information flowing in the opposite direction. In this situation, a convenient manipulation of the matrix ($T \rightarrow T^{\dagger}$) based on the ranking and reordering of the elements according to their directionality index yields an upper triangular matrix (see Materials and Methods). The transition from such hierarchical or centralized driving to a symmetric scenario can be clearly identified by monitoring the ratio $\theta = T^{\dagger}_l / T^{\dagger}_u$ between the sum of all elements of $T^{\dagger}$ in the lower triangle and the same quantity evaluated in the upper triangle. As schematically illustrated in Fig. 4, in a regime of perfectly directed driving, all the elements below the diagonal are zeros; that is, $\theta \approx 0$. In the opposite situation (that is, in the perfectly symmetric regime), the values below and above the diagonal are comparable; that is, $\theta \approx 1$. Hence, the quantity θ can be considered a suitable order parameter to characterize this order-disorder transition

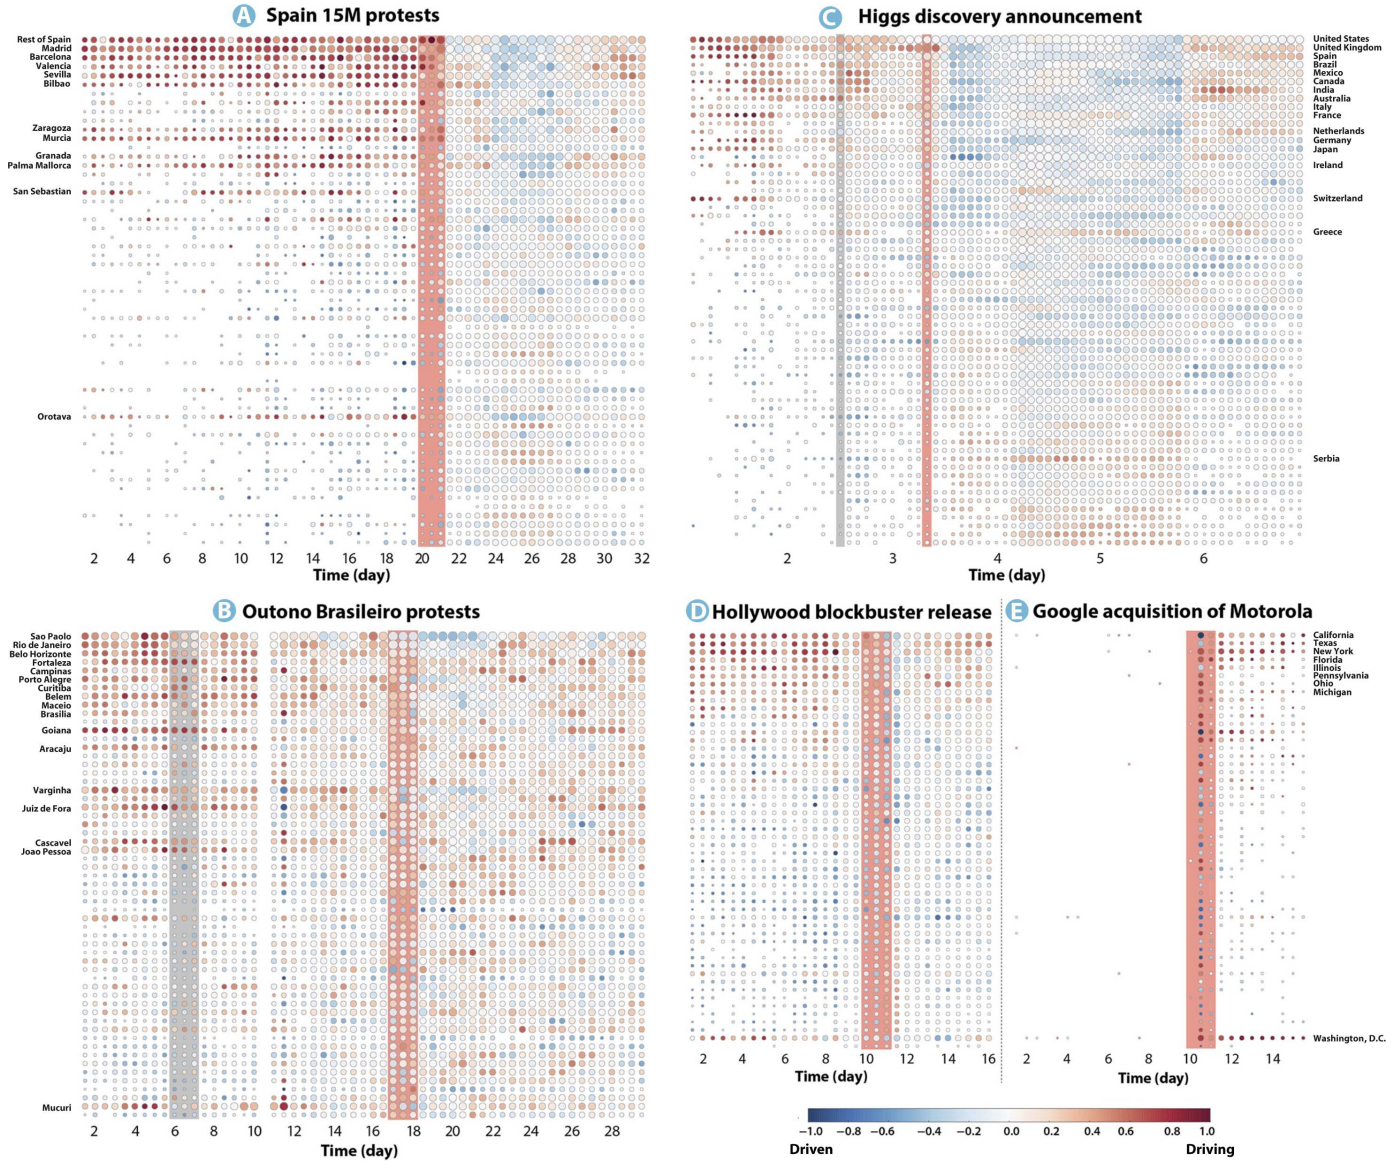

Fig. 3. Evolution of information flow balance between geographical locations for the analyzed events. (**A** to **E**) The color goes from dark blue to dark red (white corresponds to null driving), with the former standing for negative values of $\sum_Y T^S_{X,Y}$ (that is, driven locations) and the latter corresponding to positive information flow balances (that is, drivers). The size of the circles is log-proportional to the number of messages sent from the location at that time, and the vertical bars mark the day of the main event. The geographical locations are ordered according to population size, except for (C), in which countries are ranked with the amount of Higgs-related tweets produced.

and thus helps in the identification and differentiation of communication patterns across the subunits of a system.

We show in Fig. 5 the behavior of the parameter θ as a function of time in our five data sets. In all of the cases, we initially observe a highly asymmetric effective network, where a few subunits have a dominant directional coupling to the rest of the system and $\theta \ll 1$. As the systems approach the onset date of the collective event, the quantity $T_l^{\dagger}/T_u^{\dagger}$ undergoes a quick transition to $\theta \approx 1$, identifying a regime in which the couplings indicate the existence of collective phenomena where all subunits mutually affect each other. We see that in four of the five data sets, the system has a clear order-disorder transition occurring in the proximity of the collective event. In the case of the Brazilian protests, the measure significantly increases before the main event (June 17). Such behavior probably results from the effects of small precursor protests taking place from June 6 onward. The same behavior is observed in the Higgs boson data set, given that rumors started circulating after July 2. Once more, the Google data set behaves in a completely different way, never showing a clear signature of a collective regime for the couplings network. In the Supplementary Materials, we report the same analysis using the randomized signal for both the 15M and Brazil events, and we observe no order-disorder transition. Similarly, no transition exists for the Twitter unfiltered stream case study (also in the Supplementary Materials).

All data sets cover a time span preceding and following the event, and details on data collection, spatial aggregation (including keyword

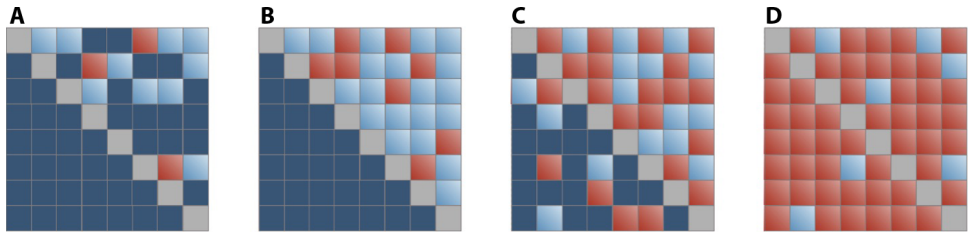

Fig. 4. Schematic representation of a transition from a centralized to a decentralized information flow scenario. If, for any given pair (x, y), $T^{S}_{x,y} \sim T_{x,y}$, all existent dynamical driving is net driving; that is, subsystems present a highly hierarchical structure. In this scenario, if a subsystem dominates another one, the former is not dominated by the latter. This is well illustrated in (**A**) and (**B**). Note, however, that only a few subsystems play an active (dynamical) role in (A), whereas the situation has reached a perfectly hierarchical structure in (B). Indeed, in this idealized situation, the net transfer entropy reaches its maximum: any further addition in terms of dynamical driving will decrease the amount of net transfer entropy [as in (**C**)]. Furthermore, (B) and (C) illustrate that there exists a tipping point beyond which the event has necessarily gone global. The extreme case where every subsystem exerts some amount of dynamical driving results in a "null driving" scenario [as in (**D**)]. In this schematic representation, the color scales go from dark blue to red, that is, zero to maximum transfer entropy, respectively.

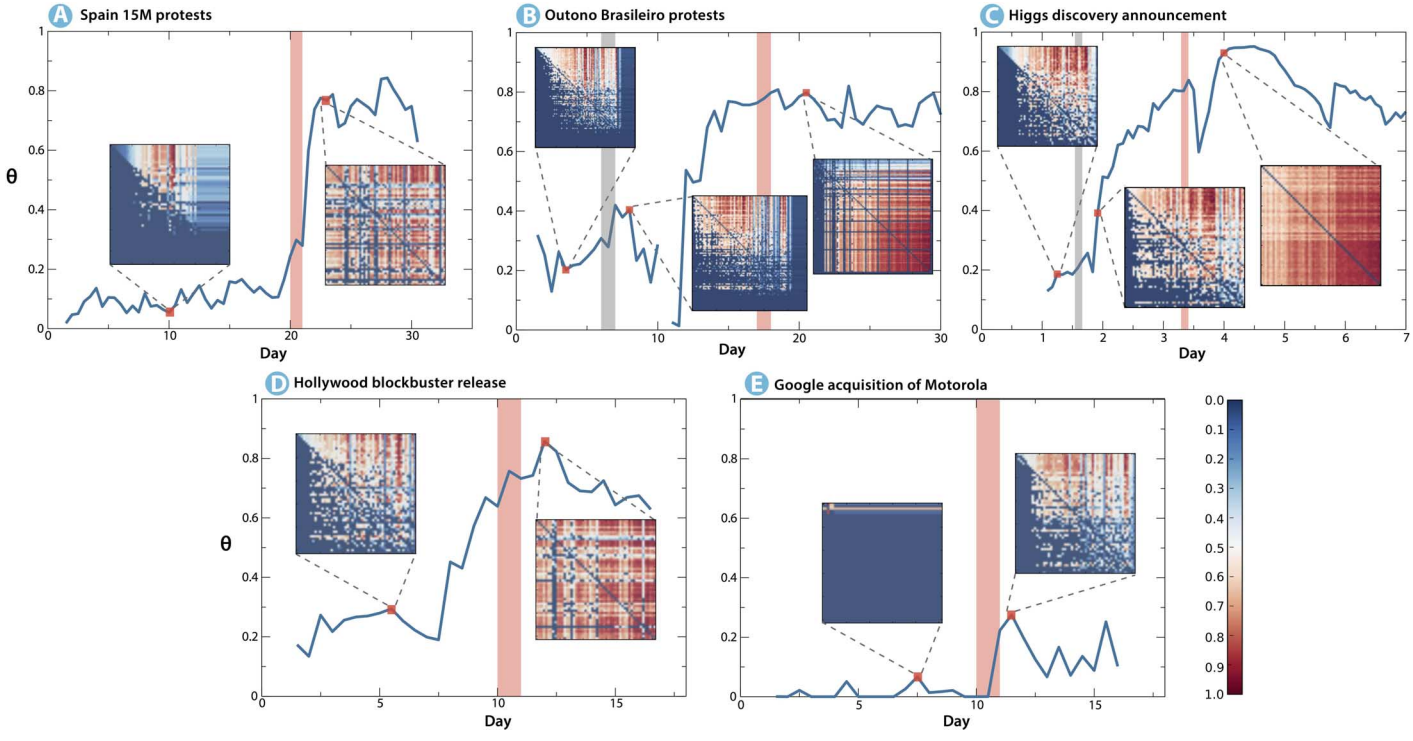

Fig. 5. Order parameter θ as a function of time for the five events analyzed. The figure represents the behavior of the ratio $\theta = T^{\dagger}_{l}/T^{\dagger}_{u}$ characterizing the order/disorder of the effective connectivity matrix as a function of time (note a point missing in the Brazilian data set because of a data blackout between days 10 and 11). For each data set, two or three matrices $T^{\dagger}$ are plotted, considering one or two times before and one after the main event (signaled with a red vertical bar). A clear transition from a hierarchical directed to a distributed symmetrical scenario is observed for events (**A**), (**B**), (**C**), and (**D**). The Google data set, depicted in (**E**), behaves differently by not showing the same evidence of transition effects.

selection and the geolocalization of messages), and sensitivity analysis of the methodology can be found in Materials and Methods and in the Supplementary Materials.

DISCUSSION

The mapping of influence networks using an information theoretical approach offers a new lens to analyze the emergence of collective phenomena. Through this lens, we have revealed the existence of a double transition—in the time scales (slow to fast) and directional couplings (hierarchical to distributed)—in systems that gather around some sort of collective action. Regarding the first, we bring to light that time series analysis should pay attention to the time scales of the underlying dynamical processes if it is to provide a reliable account of those dynamics—a fact that resonates beyond societal analysis. We also uncover the effective network of information flow between spatially defined subunits of the social system and study the structural changes of the network connectivity pattern as the system goes through different collective states. In addition, the effective network lends itself to further

analysis that can lead to the identification of structural hubs, coordinated communities, influence pathways of geographical or cultural characteristics, and geographical subunits that may have recurrent roles in the onset of social phenomena. The methodology we present here can therefore be used to gain new insights on the structural and functional relations occurring in large-scale structured populations, eventually leading to the identification of metrics that might be used for the definition of precursors of large-scale social events.

Additionally, the methodology presented here opens interesting paths to advance in the analysis of social phenomena and the identification of generative mechanisms; however, this advance should not be conflated with the possibility of forecasting the emergence of social events, but as a natural complement of techniques along this line, such as network change-point detection (*30*). The evidence we discuss is agnostic with regard to the predictive potential of online networks and microblogging platforms. A real predictive approach cannot be disentangled from an automatic selection of the relevant discussion topics. Our analyses use data sets that were already zooming into the right conversation domain and monitoring specific keywords/hashtags in the Twitter stream. However, we believe that the general methodological framework we put forward is a first step toward a better understanding of the temporal and spatial signatures of large-scale social events. This advancement might eventually inform the development of tools that can help us anticipate the emergence of macroscopic phenomena. In the meantime, our method offers a valuable resource to analyze how information-driven transitions unfold in socially relevant contexts.

MATERIALS AND METHODS

Data

The first data set focuses on the Spanish 15M movement, which emerged in 2011 (*31*, *32*). The data cover a dormant period of low microblogging activity that is followed by an explosive phase in which the movement gained the attention of the general public and was widely covered by traditional media sources (see Fig. 1). The second data set contains more than 2.5 million geolocalized tweets associated to the Outono Brasileiro, a set of political protests that emerged in Brazil in June 2013. Similar to the Spanish case, the Brazilian data include an initial phase of low activity followed by a gradual escalation toward the high volumes of general attention that accompanied the street protests. The third data set tracks communication on the discovery of the Higgs boson before and after it was officially announced to the press in July 2012; this data set has been used before to assess how rumors spread through online social networks (*29*). The fourth data set contains messages related to the release of a Hollywood blockbuster, announced months before its premiere to stir momentum among the fan base. Finally, we also consider a data set that tracks communication on the acquisition of Motorola by Google, which came as sudden and unexpected news and immediately triggered a high volume of public attention.

Spanish Twitter activity is spatially coarse-grained according to the list of metropolitan areas defined by the European Spatial Planning Observation Network (*33*). This process yields 56 aggregated time series, each of which corresponds to a different geographical area. In addition, there is an extra signal that accounts for any activity not included in those areas; that is, the system is made up of $n = 57$ components. The data from Brazil are aggregated in 97 basins, which correspond roughly to metropolitan areas (*34*, *35*). The data that track rumors about the Higgs boson are aggregated at the country level, including only the $n = 61$ most active basins around this topic. Finally, the Motorola-Google and the blockbuster data are classified in 52 U.S. areas: 50 states plus Washington, DC, and Puerto Rico.

Order-disorder transition

In real data sets, the transition between centralized and decentralized information-transfer scenarios can be visually inspected with a convenient sorting of the rows and columns of the $T_{x,y}$ matrix. We did so in Fig. 5, ranking each subunit of the system. The rank for a subunit x is assigned according to the number of times x is dominant over the rest of the subunits. Once the ranking is settled, any $T_{x,y} < \frac{1}{2} T^{\max}_{x,y}$ is set to 0 to improve the visual understanding of the figure. We then obtained a transformed matrix, that is, $T_{x,y} \rightarrow T^{\dagger}_{x,y}$. Beyond visualization, the sorted matrix allows us to define a monitoring measure $\theta = \frac{\sum_{x>y} T^{\dagger}_{x,y}}{\sum_{x<y} T^{\dagger}_{x,y}} = \frac{T^{\dagger}_{\mathrm{l}}}{T^{\dagger}_{\mathrm{u}}}$ (that is, the ratio between the sums of all the matrix's elements in the lower and upper triangles), which provides a quantification of the state of the system (as explained in the main text and illustrated in Fig. 4). For completion, we have also plotted the same figures without threshold (see section B.6 in the Supplementary Materials).

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at http://advances.sciencemag.org/cgi/content/full/2/4/e1501158/DC1

Data, context, and chronology of the events analyzed

Methods used in the analysis

Sensibility analysis of the parametrization

Validation of results (I): Time series randomization

Validation of results (II): Unfiltered Twitter stream

Validation of results (III): Synthetic time series generation

Table S1. List of keywords used to find tweets related to the Outono Brasileiro.

Fig. S1. Schematic representation of the algorithm used to gather geographical coordinates of the Hollywood movie release and the Google-Motorola acquisition data sets.

Fig. S2. Sample order pattern for $m = 3$.

Fig. S3. Schematic view of the sliding window scheme.

Fig. S4. Evolution of the order parameter θ for thresholded (green) and raw (red) $T^{\dagger}$ matrices.

Fig. S5. Dependence of τ with the sliding window size ω, considering the Spanish 15M protest.

Fig. S6. Normalized directionality index for each geographical unit in the 15M data set for different ω.

Fig. S7. Fraction of false nearest neighbors as a function of m for the Spanish data set and the Madrid time series.

Fig. S8. Normalized directionality index for each geographical unit in the 15M data set for different m.

Fig. S9. Characteristic time scale τ for four data sets at alternative geographical aggregation levels.

Fig. S10. Normalized directionality index for four data sets at alternative geographical aggregation levels.

Fig. S11. Behavior of θ as a function of time for four data sets at alternative geographical aggregation levels.

Fig. S12. Average total amount of STE for some Δt (top panel) and time scale profile τ (bottom panel) for 15M data set amplitude adjusted Fourier transform surrogates (50 randomizations).

Fig. S13. Behavior of θ as a function of time for 15M and Outono Brasileiro data sets randomized surrogates.

Fig. S14. Average total amount of STE for some Δt (top panel) and time scale profile τ (bottom panel) for 15M data set constrained surrogates (20 randomizations).

Fig. S15. Evolution of τ as a function of time.

Fig. S16. Thresholded $T^{\dagger}$ matrices corresponding to different moments in the Twitter unfiltered data set.

Fig. S17. Raw time series for Twitter unfiltered stream for $\Delta t = 600$ s and $\Delta t = 45$ s (left and right, respectively).

Fig. S18. Evolution of two nonlinear systems under four changing scenarios: from dynamic independence (β = 0) to strong asymmetric coupling (β = 20.0).
Part 1. Minimalist example: Disentangling volume and time scales (Δ*t*).
Part 2. Nonlinear Lorentz oscillators: Time scales, volume, and dynamical coupling.
References (*36*–*58*)

REFERENCES AND NOTES

1. A. Culotta, *Towards Detecting Influenza Epidemics by Analyzing Twitter Messages* (ACM, New York, 2010), pp. 115–122.
2. J. Ginsberg, M. H. Mohebbi, R. S. Patel, L. Brammer, M. S. Smolinski, L. Brilliant, Detecting influenza epidemics using search engine query data. *Nature* **457**, 1012–1014 (2009).
3. K. S. Hickmann, G. Fairchild, R. Priedhorsky, N. Generous, J. M. Hyman, A. Deshpande, S. Y. Del Valle, Forecasting the 2013–2014 influenza season using Wikipedia. *PLOS Comput. Biol.* **11**, e1004239 (2015).
4. P. Chakraborty, P. Khadivi, B. Lewis, A. Mahendiran, J. Chen, P. Butler, E. O. Nsoesie, S. R. Mekaru, J. S. Brownstein, M. Marathe, N. Ramakrishnan, Forecasting a moving target: Ensemble models for ILI case count predictions, *Proceedings of the 2014 SIAM International Conference on Data Mining* (Society for Industrial and Applied Mathematics, 2014), pp. 262–270.
5. S. Asur, B. A. Huberman, *Predicting the Future with Social Media, WI-IAT '10* (IEEE Computer Society, Washington, DC, 2010), pp. 492–499.
6. J. Bollen, H. Mao, X. Zeng, Twitter mood predicts the stock market. *J. Comput. Sci.* **2**, 1–8 (2011).
7. C. Curme, T. Preis, H. E. Stanley, H. S. Moat, Quantifying the semantics of search behavior before stock market moves. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 11600–11605 (2014).
8. A. Tumasjan, T. O. Sprenger, P. G. Sandner, I. M. Welpe, Predicting elections with twitter: What 140 characters reveal about political sentiment. *ICWSM* **10**, 178–185 (2010).
9. A. Livne, M. P. Simmons, E. Adar, L. A. Adamic, The party is over here: Structure and content in the 2010 election, *Proceedings of the Fifth International Conference on Weblogs and Social Media*, Barcelona, Spain, 17 to 21 July 2011 (ICWSM).
10. F. Ciulla, D. Mocanu, A. Baronchelli, B. Gonçalves, N. Perra, A. Vespignani, Beating the news using social media: The case study of American idol. *EPJ Data Sci.* **1**, 1–11 (2012).
11. F. Botta, H. S. Moat, T. Preis, Quantifying crowd size with mobile phone and *Twitter* data. *R. Soc. Open Sci.* **2**, 150162 (2015).
12. J. Xu, T. C. Lu, R. Compton, D. Allen, in *Social Computing, Behavioral-Cultural Modeling and Prediction, Lecture Notes in Computer Science*, W. Kennedy, N. Agarwal, S. Yang, Eds. (Springer International Publishing, Cham, Switzerland, 2014) vol. 8393, pp. 403–411.
13. N. Ramakrishnan, P. Butler, S. Muthiah, N. Self, R. Khandpur, P. Saraf, W. Wang, J. Cadena, A. Vullikanti, G. Korkmaz, C. Kuhlman, A. Marathe, L. Zhao, T. Hua, F. Chen, C.-T. Lu, B. Huang, A. Srinivasan, K. Trinh, L. Getoor, G. Katz, A. Doyle, C. Ackermann, I. Zavorin, J. Ford, K. Summers, Y. Fayed, J. Arredondo, D. Gupta, D. Mares, 'Beating the News' with EMBERS: Forecasting Civil Unrest Using Open Source Indicators, KDD '14 (ACM, New York, 2014), pp. 1799–1808.
14. M. Skoric, N. Poor, P. Achananuparp, E. P. Lim, J. Jiang, Tweets and votes: A study of the 2011 Singapore General Election, *IEEE in 2012 45th Hawaii International Conference on System Sciences* (2012), pp. 2583–2591.
15. E. T. K. Sang, J. Bos, Predicting the 2011 Dutch senate election results with Twitter, *Proceedings of SASN 2012, the EACL Workshop on Semantic Analysis in Social Networks (ACL)*, Avignon, France, 2012.
16. D. Gayo-Avello, "I wanted to predict elections with twitter and all I got was this lousy paper"—A balanced survey on election prediction using Twitter data. arXiv:1204.6441 (2012).
17. Z. Tufekci, Big questions for social media big data: Representativeness, validity and other methodological pitfalls. arXiv:1403.7400 (2012).
18. D. M. Lazer, R. Kennedy, G. King, A. Vespignani, Big data. The parable of Google Flu: Traps in big data analysis. *Science* **343**, 1203–1205 (2014).
19. D. Helbing, Globally networked risks and how to respond. *Nature* **497**, 51–59 (2013)
20. J. Lehmann, B. Goncalves, J. J. Ramasco, C. Cattuto, Dynamical classes of collective attention in twitter, *Proceedings of the 21st International Conference on World Wide Web (WWW'12)*, ACM, New York, 2012.
21. M. Staniek, K. Lehnertz, Symbolic transfer entropy. *Phys. Rev. Lett.* **100**, 158101 (2008).
22. C. Bandt, B. Pompe, Permutation entropy: A natural complexity measure for time series. *Phys. Rev. Lett.* **88**, 174102 (2002).
23. T. Schreiber, Measuring information transfer. *Phys. Rev. Lett.* **85**, 461–464 (2000).
24. K. Hlaváčková-Schindler, M. Paluš, M. Vejmelka, J. Bhattacharya, Causality detection based on information-theoretic approaches in time series analysis. *Phys. Rep.* **441**, 1–46 (2007).
25. K.-Y. Ni, T.-C. Lu, Information dynamic spectrum characterizes system instability toward critical transitions. *EPJ Data Sci.* **3**, 28 (2014).
26. J. T. Lizier, J. Heinzle, A. Horstmann, J.-D. Haynes, M. Prokopenko, Multivariate information-theoretic measures reveal directed information structure and task relevant changes in fMRI connectivity. *J. Comput. Neurosci.* **30**, 85–107 (2011).
27. S. Kullback, R. A. Leibler, On information and sufficiency. *Ann. Math. Stat.* **22**, 79–86 (1951).
28. O. Sporns, D. R. Chialvo, M. Kaiser, C. C. Hilgetag, Organization, development and function of complex brain networks. *Trends Cogn. Sci.* **8**, 418–425 (2004).
29. M. De Domenico, A. Lima, P. Mougel, M. Musolesi, The anatomy of a scientific rumor. *Sci. Rep.* **3**, 2980 (2013).
30. L. Peel, A. Clauset, Detecting change points in the large-scale structure of evolving networks, *Proceedings of the 29th International Conference on Artificial Intelligence (AAAI)*, pp. 2914–2920 (2015).
31. J. Borge-Holthoefer, A. Rivero, I. García, E. Cauhé, A. Ferrer, D. Ferrer, D. Francos, D. Iñiguez, M. P. Pérez, G. Ruiz, F. Sanz, F. Serrano, C. Viñas, A. Tarancón, Y. Moreno, Structural and dynamical patterns on online social networks: The Spanish May 15th movement as a case study. *PLOS One* **6**, e23883 (2011).
32. S. González-Bailón, J. Borge-Holthoefer, A. Rivero, Y. Moreno, The dynamics of protest recruitment through an online network. *Sci. Rep.* **1**, 197 (2011).
33. European Spatial Planning Observation Network, www.espon.eu [accessed 16 April 2014].
34. D. Balcan, B. Gonçalves, H. Hu, J. J. Ramasco, V. Colizza, A. Vespignani, Modeling the spatial spread of infectious diseases: The GLobal Epidemic and Mobility computational model. *J. Comput. Sci.* **1**, 132–145 (2010).
35. D. Balcan, H. Hu, B. Goncalves, P. Bajardi, C. Poletto, J. J. Ramasco, D. Paolotti, N. Perra, M. Tizzoni, W. Van den Broeck, V. Colizza, A. Vespignani, Seasonal transmission potential and activity peaks of the new influenza A(H1N1): A Monte Carlo likelihood analysis based on human mobility. *BMC Med.* **7**, 45 (2009).
36. American regions, http://en.wikipedia.org/wiki/List_of_regions_of_the_United_States [accessed 8 August 2015].
37. Brazilian states, http://en.wikipedia.org/wiki/States_of_Brazil[accessed 8 August 2015].
38. Spain's autonomous communities, http://en.wikipedia.org/wiki/Autonomous_communities_of_Spain [accessed 8 August 2015].
39. D. Balcan, V. Colizza, B. Gonçalves, H. Hu, J. J. Ramasco, A. Vespignani, Multiscale mobility networks and the large scale spreading of infectious diseases. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 21484–21489 (2009).
40. M. Castells, *Networks of Outrage and Hope: Social Movements in the Internet Age* (John Wiley and Sons, New York, 2013).
41. *Center for International Earth Science Information Network (CIESIN) and Centro Internacional de Agricultura Tropical (CIAT)*. The Gridded Population of the World Version 3 (GPWv3): Population Grids (Socioeconomic Data and Applications Center, Columbia University, Palisades, New York), http://sedac.ciesin.columbia.edu/data/collection/gpw-v3.
42. *International Food Policy Research Institute (IF- PRI), The World Bank, Center for International Earth Science Information Network (CIESIN) and Centro Internacional de Agricultura Tropical (CIAT)*. Global Rural-Urban Mapping Project (GRUMP), Alpha Version: Population Grids (Socioeconomic Data and Applications Center, Columbia University, Palisades, New York), http://sedac.ciesin.columbia.edu/data/collection/gpw-v3.
43. M. D. Conover, C. Davis, E. Ferrara, K. McKelvey, F. Menczer, A. Flammini, The geospatial characteristics of a social movement communication network. *PLOS One* **8**, e55957 (2013).
44. GeoNames, www.geonames.org/ [accessed 2012].
45. P. Gerbaudo, Tweets and the streets: Social media and contemporary activism (Pluto Press, London, 2012).
46. C. W. J. Granger, Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* **37**, 424–438 (1969).
47. J. J. Jones, R. M. Bond, C. J. Fariss, J. E. Settle, A. D. I. Kramer, C. Marlow, J. H. Fowler, Yahtzee: An anonymized group level matching procedure. *PLOS One* **8**, e55760 (2013).
48. M. B. Kennel, R. Brown, H. D. I. Abarbanel, Determining embedding dimension for phase-space reconstruction using a geometrical construction. *Phys. Rev. A* **45**, 3403–3411 (1992).
49. S. Schinkel, N. Marwan, J. Kurths, Order patterns recurrence plots in the analysis of ERP data. *Cogn. Neurodyn.* **1**, 317–325 (2007).
50. T. Schreiber, TISEAN software; www.mpipks-dresden.mpg.de/~tisean/Tisean_3.0.1.
51. T. Schreiber, Constrained randomization of time series data. *Phys. Rev. Lett.* **80**, 2105–2108 (1998).
52. T. Schreiber, A. Schmitz, Improved surrogate data for nonlinearity tests. *Phys. Rev. Lett.* **77**, 635–638 (1996).
53. C. E. Shannon, A mathematical theory of communication, 1948. *Bell Syst. Tech. J.* **27**, 379–423, 623–656 (1948).
54. S. A. Soule, The student divestment movement in the United States and tactical diffusion: The shantytown protest. *Social Forces* **75**, 855–882 (1997).
55. K. T. Andrews, M. Biggs, The dynamics of protest diffusion: Movement organisations, social networks, and news media in the 1960 sit-ins. *Am. Sociol. Rev.* **71** 752–777 (2006).
56. R. K. Givan, K. M. Roberts, S. A. Soule, Eds., *The Diffusion of Social Movements: Actors, Mechanisms, and Political Effects* (Cambridge Univ. Press, Cambridge, 2010).
57. D. J. Wang, S. A. Soule, Social movement organizational collaboration: Networks of learning and the diffusion of protest tactics, 1960–1995. *Am. J. Sociol.* **117**, 1674–1722 (2012).
58. B. Carbunar, R. Potharaju, You unlocked the Mt. Everest badge on Foursquare! Countering location fraud in geosocial networks, *IEEE 9th International Conference on Mobile Adhoc and Sensor Systems (MASS)*, pp. 182–190 (2012).

Acknowledgments: We thank D. Allen, R. Compton, and T.-C. Lu at HRL Laboratories LLC for assistance with the Brazilian data set and useful discussions; we also thank A. Lima for sharing the Higgs boson data. **Funding:** B.G. thanks the Moore and Sloan Foundations for support as part of the Moore-Sloan Data Science Environment at New York University. A.A. acknowledges the support of the European Union MULTIPLEX (317532) and the Spanish Ministry of Science and Innovation (FIS2012-38266-C02-01) and partial financial support from the ICREA Academia and the James S. McDonnell Foundation. Y.M. acknowledges support from MINECO through grant FIS2011-25167, Comunidad de Aragón (Spain) through a grant to the group FENOL, and the European Commission Future and Emerging Technologies Proactive Project MULTIPLEX through grant 317532. For the analysis of data outside of the United States, A.V. and N.P. acknowledge the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center (DoI/NBC) contract number D12PC00285. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBE, or the United States Government. The funders had no role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript. **Author contributions:** All authors contributed equally to this work. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper, the Supplementary Materials, and/or www.jbh.cat/data/twitter_timeseries.tar.gz. Additional data related to this paper may be requested from the authors.

Submitted 25 August 2015
Accepted 18 February 2016
Published 1 April 2016
10.1126/sciadv.1501158

Citation: J. Borge-Holthoefer, N. Perra, B. Gonçalves, S. González-Bailón, A. Arenas, Y. Moreno, A. Vespignani, The dynamics of information-driven coordination phenomena: A transfer entropy analysis. *Sci. Adv.* **2**, e1501158 (2016).

ScienceAdvances
AAAS

The dynamics of information-driven coordination phenomena: A transfer entropy analysis

Javier Borge-Holthoefer, Nicola Perra, Bruno Gonçalves, Sandra González-Bailón, Alex Arenas, Yamir Moreno and Alessandro Vespignani (April 1, 2016)

Sci Adv 2016, 2:.
doi: 10.1126/sciadv.1501158

Science Advances

AAAS

advances.sciencemag.org/cgi/content/full/2/4/e1501158/DC1

Supplementary Materials for

The dynamics of information-driven coordination phenomena: A transfer entropy analysis

Javier Borge-Holthoefer, Nicola Perra, Bruno Gonçalves, Sandra González-Bailón, Alex Arenas, Yamir Moreno, Alessandro Vespignani

Published 1 April 2016, *Sci. Adv.* **2**, e1501158 (2016)
DOI: 10.1126/sciadv.1501158

The PDF file includes:

A. Data, Context, and Chronology of the Events Analyzed

We considered five different events: the Spanish 15M protests, the *Outono Brasileiro* (Brazilian autumn) movement, the announcement of the Higgs boson discovery, the release of a Hollywood blockbuster movie (Batman "The Dark Knight Rises"), and the acquisition of Motorola by Google. In this section we report details concerning these events and the associated Twitter data sets.

A.1. The 15M Protests (May 2011).

These protests emerged in Spain in the aftermath of the so-called Arab Spring. A grassroots social movement, later called the *Indignados* ("the outraged"), it emerged from online communication amongst a decentralized network of citizens and civic associations. Online networks (blogs, Facebook, Twitter) were used to spread a call for action for May 15, 2011. The main drivers of the protests were spending cuts and policy reactions to the economic crisis. Massive demonstrations took place on May 15 in several major cities around Spain, many of them resulting in camp sites in main city squares that remained active for weeks. Mainstream media didn't cover the movement until it reached the streets. As a consequence, most communication and broadcasting announcing and discussing the mobilizations took place through online channels. Social media networks (in particular, Twitter) served a crucial role in the coordination of the protests and the management of camp logistics [40,45].

The Twitter data for the Spanish 15M movement were harvested by a startup company (*Cierzo Ltd.*) for a period spanning from April 25 to May 25, 2011. The main demonstrations in Spain took place on May 15th and onwards, thus our analysis covers a brewing period with low activity rates (up to May 15th, day 20 in the Figures) plus an "explosive" phase beyond that date, in which the phenomenon reached general public and was widely covered by traditional mass media, see Figure 1 in the Main Text. Scraps on Twitter servers yielded 581,749 messages.

A.2. Outono Brasileiro Protests (June 2013)

More recently, massive protests filled the streets of several Brazilian cities. The triggering factor was the rising prices of public transportation, but on the background loomed long-standing discontent with inequality, the government economic policies, and the provision of social services. Social media played again an instrumental role in the coordination of large-scale mobilization and the swift diffusion of information: images documenting the often brutal police reaction to the protests boosted mobilization and brought more people to the streets of more cities and municipalities. The protests, often dubbed as *Outono Brasileiro* ("Brazilian autumn"), resulted in Brazilian President Dilma Rouseff announcing, in June 21, measures to improve the management of public transport along with other social services. This prime-time televised address, however, did not placate citizens dissatisfaction, who continued staging protests in subsequent days.

The data set regarding such event has been obtained using the *PowerTrack* tool that provides 100% coverage for a set of specified keywords (see Table S1). For our analysis we considered just the tweets sent in the month of June 2013 (2,670,933 tweets). Indeed, the first large scale protest, often associated with the escalation of the protests, took place on June 17th, with remarkable (though smaller) precursors on the 6th and 13th. As in the case of the Spanish movement we considered a brewing period with low activity rates plus the "explosive" phase beyond the date of the first massive street protest.

AnonymousBrasil	boicot	cacerolada	cacerolazo	huelga
marcha	marchado	marcham	marchamos	marche
march\'e	marcho	Passeata	protesta	protestam
protestar\'as	protestarem	protestarmos	proteste	protestemos
protesten	protesto	protest\'o	concentraci\'on	reforma
greve	rali	manifestaçao	manifestantes	corrup\c c\~ao

Table S1. List of keywords used to find tweets related to the Outono Brasileiro.

A.3. The Higgs Boson Discovery Announcement (July 2012)

In July 4 2012, a team of scientists based at CERN presented results that indicated the existence of a new particle, compatible with the Higgs boson (the existence of which had first been hypothesized in 1964). Mainstream news media covered the discovery after the announcement, but during the days preceding it there were already rumors of its discovery circulating through social media [29]. The messages we analyze were collected using Twitter's publicly available API between July 1st and July 7th using a list of relevant keywords (i.e. LHC, CERN, boson, Higgs). In total, the data set contains 985,590 tweets.

A.4. The Hollywood Movie Release (July 2012)

The Dark Knight Rises is the third installment of the Batman trilogy (started in 2005 with the release of Batman Begins and followed up in 2008 with The Dark Knight). It was premiered in New York on July 16th 2012, and released in several English-speaking countries a few days later. The promotional campaign included so-called viral marketing through social media. The film was nominated to several prestigious awards, and grossed over a billion dollars in the box office.

The data set includes 130,529 tweets between July 6th and July 21st that include the words “batman”, “darkknight” or “darkknightrises”. The tweets are obtained from the Twitter Gardenhose (a 10% random sample of the entire Twitter traffic).

A.5. The Google-Motorola Acquisition (August 2011)

On August 15th 2011, Google announced a relatively unexpected agreement to acquire the mobile company Motorola. The move was a strategic attempt to strengthen Google’s patent portfolio in a context where legal battles over patents is increasingly shaping the mobile industry and the telecommunications environment.

The data set contains 10,890 tweets between August 5th and August 20th, 2011. In order to minimize the noise, we considered just tweets containing both “Google” and “Motorola”. Also in this case, the tweets are obtained from the Twitter Gardenhose.

B. Methods used in the analysis

In this section we detail all the methods that have been used throughout the work: Twitter time series construction and temporal coarse-graining (B.1, B.2); Symbolic Transfer Entropy (STE) measurement (B.3); time scale determination as a maximization of STE (B.4); and central-to-distributed transition characterization (B.5, B.6).

B.1. Data spatial aggregation

With activity information at hand, a possible way to represent information is to assign a time series to individual Twitter users. This however has important drawbacks: activity may be too sparse to build a significant series; it may be rather difficult to detect general, meaningful trends when studying series interaction; finally, one needs to take into account computational costs. Thus, we have chosen to coarse-grain the data from a geographical point of view. Some considerations need to be made at this point:

I. Although the STE does not consider signal volume, the geographical partition has an upper and lower cut-off for the information significance of the social media signal. Generally speaking, if we go down to the scale of a single village or neighborhood, alongside computational issues coming from the dimensionality of the data, driver and driven regions will be too fragmented to provide a clear pattern. On the contrary, if we aggregate the signal on too large regions we will be integrating all partitions into a unique driver area thus making the signal homogeneous and uninformative. This is a long standing issue that affects several other areas ranging from gravity law in transportation science to epidemic modeling in meta-population structure. In all these cases, as in our approach, a best practice is to perform sensitivity analysis for the identification of a meaningful geographic scale.

II. Other alternative partitions are possible: instead of geographical boundaries, we could have used age, religion, gender or other grouping categories, if that information had been available. It is likely that alternative partitions yield even higher STE values, placing a challenging multi-optimization problem. In this work, the quest for an STE-optimal partition is beyond the scope: we chose events for which geography makes sense, i.e. a large corpus of sociological literature has highlighted in the past the role of geography in information diffusion, economic transactions, personal interactions, and civil and political movements [43,47, 54-57]

For the 15M case, geographical information was collected for each user involved in the protests, thereafter tweets were assigned their author's location. Spanish Twitter activity is spatially coarse-grained according to the list of metropolitan areas defined by the European Spatial Planning Observation Network (http://www.espon.eu). This process yields 56 aggregate time series, each corresponding to a geographical area, plus an extra signal which accounts for any activity not included in the previous definition. Thus, the system is made up of N=57 components. Time-stamps have been modified when necessary (Santa Cruz de Tenerife, Orotava and Palmas de Gran Canaria) to a common time frame. The pre-defined metropolitan areas account for over half Spain's total population.

The Brazilian tweets have been instead aggregated at the level of N=97 basins centered around major transportation hubs. These geographical units, that correspond to census areas surrounding large cities, have been defined aggregating population cells of 15×15 minutes of arc (approximately a rectangle of 25×25 kms along the Equator), from the "Gridded Population of the World" and the "Global Urban-Rural Mapping" projects [41,42], to the closest airport that satisfies the following two conditions: (i) Each cell is assigned to the closest airport within the same country; and (ii) the distance between the

airport and the cell cannot be longer than 200km. This cutoff naturally emerges from the distribution of distances between cells and closest airports. See refs. [35,39] for details. Moreover, having access to 100% of the entire signal on Twitter associated with at least one word listed in Table S1, we considered just tweets with live GPS coordinates.

Tweets around the Higgs boson discovery were aggregated at the country level. The original data set contained tweets from over 200 countries, but these have been thresholded to retain only those countries with more than 500 tweets over the topic of interest, for the whole week. This entails that only N=61 countries are present in the analysis in the main text. The details regarding the location technique can be found in [29].

The remaining two data sets have been aggregated in N=52 areas –50 U.S. states, plus Washington D.C and Puerto Rico. In particular, the geographical information of tweets in this case has been gathered either from live GPS locations, or mining the so called "self-reported location". In general this field is filled freely by the users that can report their location at different levels. i.e. NYC, California, CA, USA etc.. Some fraction of the reported locations are jokes, i.e. Moon, Mars, behind you, etc. We parsed these fields trying to match a country, state, or city name. In fig. S1 we report the flow chart of the algorithm used.

Interestingly, the method is able to find a match of 40-50% the total number of tweets. It must be noted that identifying the location of a user is a hard problem, so some assumptions must be made, and with each assumption noise is introduced in the data. Even the simplest assumption possible, that a user is where her GPS coordinates say she is has been shown to not always be true [58]. It's also clear that the user is not always where she claims to be on her self-reported location field, as this field is infrequently updated even when users travel. On the other hand, an argument can be made that the self-reported location field (even if it does not correspond to the instantaneous location of a user) corresponds to the city or region that the user identifies with, which is relevant for our case.

In summary, while we acknowledge that the two types of location may correspond to two different things (current location *vs.* "home" location), these are precisely the two types of location that are relevant to the aim of identifying which users are invested in the political activities occurring in each city, as (in general) we are directly affected both by what is happening around us and by recent developments back home while away. If our goal is to understand how developments in one city are driving what is happening in another city, we must take the most significant location information available. We believe that both a GPS coordinate and self reported location are relevant.

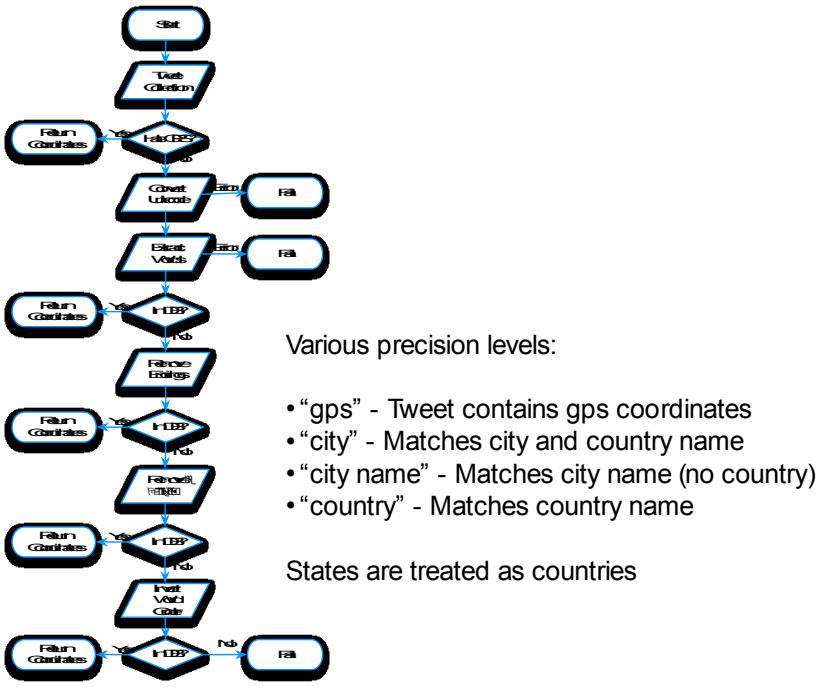

fig. S1. Schematic representation of the algorithm used to gather geographical coordinates of the Hollywood movie release and the Google-Motorola acquisition data sets. In the figure DB stands for database. We used the "GeoNames" database [44].

In subsection C.3 (below) we show the results obtained considering different geographical aggregations. Namely, we also considered N=16 Spanish communities (one of them was left out because no messages were collected from there), N=27 Brazilian states, and N=9 regions in the USA defined by the American Census Bureau [36].

B.2. Data temporal aggregation

The definition of the temporal aggregation of Twitter data is particularly important in our approach. Indeed, we want to determine the characteristic time scale at which the driving between series is most evident. The data comes with temporal resolution down to a second. However, such level of resolution is excessive to detect dynamical trends among series. We considered different sampling rates Δt spanning from 1 to 120 minutes. Although arbitrary, this range of temporal aggregations account for the fluidity of Twitter's discussion as well as the limited attention time span of users. In subsection B.4 we discuss the ideas that allow us to determine the optimal temporal aggregation scheme.

B.3. Symbolic Transfer Entropy

Closely related to other measures, such as mutual information [53], Granger causality [46] and transfer entropy [23], Symbolic Transfer Entropy (STE) [21] provides a solid method to detect and quantify the strength and direction of couplings between components of dynamical systems. The symbolic approach, on the other hand, links STE to order patterns and symbolic dynamics [22,49] as a means to successfully analyze time series which may be noisy, short and/or non-stationary.

Once spatial and temporal aggregation schemes are fixed, we proceed to measure STE as a way to quantify the coupling among series. Note that such series span long times, L, of several days or even a month. Also, activity during these days is changing due to offline events happening outside the Twittersphere. Thus, STE is not measured over time series taken as a whole, but over sliding windows of length $\omega << L$ (which is indeed a standard way to proceed in other relevant fields such as neuroscience). To obtain a finer analysis, these windows advance at a slow pace of only 30 minutes. In practice, this means that the first window spans the interval $[0,\omega]$; the second one $[30, \omega + 30]$ (in minutes), and so on. Window width ω, admittedly, is the first parameter that will affect the measurements output, and we will discuss its effects later (subsection C.1 below).

Given a window of width ω, the resulting series are transformed into symbol sequences as described in [21], for which an embedding dimension $3 \leq m \leq 7$ [22] must be chosen (see also subsection C.2 for further details). Let us consider a simple example of how this works. Imagine we have a signal

$$x = \{120, 74, 203, 167, 92, 148, 174, 47\}$$

(let us ignore sliding windows by now). We shall transform this series into symbol series. For simplicity, let us suppose that the embedding dimension m=3. This quantity determines the amount of symbols that can possibly exist, which is $m! = 6$ in our case. See fig. S2 as an illustration of the possible symbols that can be obtained.

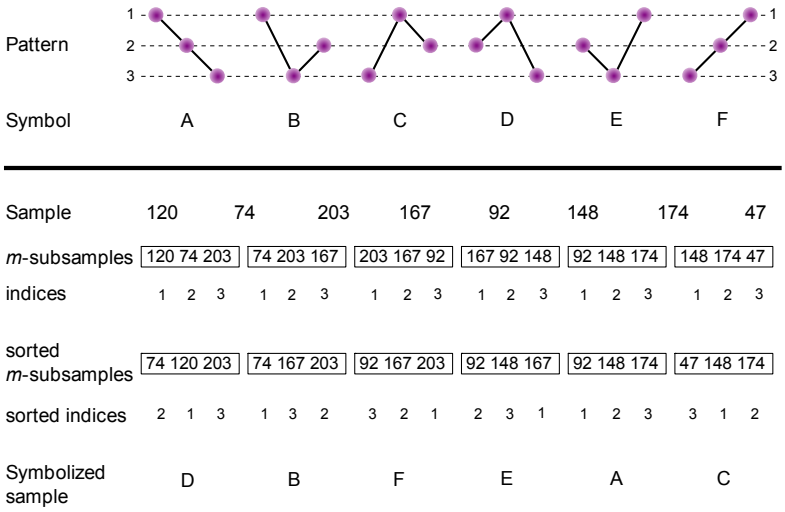

fig. S2. Sample order pattern for $m = 3$. If we neglect ties, the number of possible patterns is 3! (see [21] to check how ties are dealt with).}

The first step to transform x into symbol sequences is to sort their subchains of length m in increasing order. So, we take the first three elements of x and sort them, which leaves us with {74, 120, 203}. We have kept track of these values' indices, such that the sequence now looks like {2, 1, 3}. According to fig. S2 (top), this first subchain maps to the symbol D.

From this scheme, we just need to advance one value at a time: the next subchain to consider is {74, 203, 167}. Its sorted version is {74, 167, 203}, which corresponds to {1, 3, 2}, and maps to B. The whole process for the x signal looks like fig. S2 (center), and their sorted indices lead to fig. S2 (bottom), rendering a symbol sequence $\hat{x} = \{D, B, F, E, A, C\}$. With a similar procedure, other series y are transformed into $\hat{y}$. Given these symbol sequences $\hat{x}$ and $\hat{y}$, STE between a pair of signals (x, y) is defined as in Equation (1) of the Main Text.

A few facts need to be highlighted at this point:

I. A signal with an original length of n points is reduced, through symbolization, to a new string with n-m+1 symbols.

II. A way to interpret the meaning of m is to think of it as the amount of "expressiveness" it allows to the original series. That is, if m is low, a rich signal (one with many changes in it) is reduced to a small amount of possible symbols. This is of utmost importance to understand why we have chosen a relatively high m to work with (see subsection C.2).

III. $m = 7$ is not a *theoretical* upper limit, but rather a *practical* one: as the alphabet size scales as $m!$ the computational costs beyond that value become unfeasible (we follow ref. 22 in this aspect). Besides, there exist ways to quantify how suitable a certain embedding dimension is, i.e. false nearest neighbors (fig. S7 below). In most cases, $m > 5$ does not imply a better projection of the data in the symbol space.

IV. All measurements in the present work have been performed using $\delta = 1$. This implies that we are measuring the capacity of a signal to predict the *immediate* future of another signal, i.e. just one symbol ahead.

Finally, the following figure summarizes how STE has been measured in all data sets throughout the main text and SM, specifying the meaning of each parameter (s, m, Δt, etc.) and the range of values that have been used in each case.

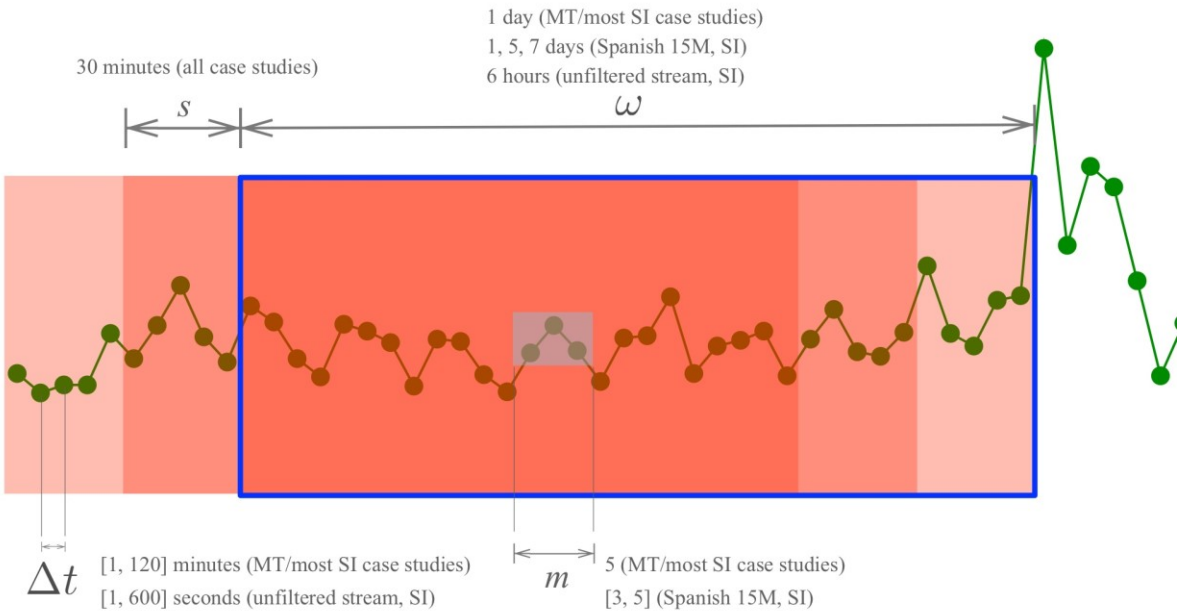

fig. S3. Schematic view of the sliding window scheme. Each snapshot of width ω encompasses a sub-chain of ω / Δ*t* points. This sub-chain is symbolized, taking *m* time series ticks to construct a single symbol. Once STE has been measure for a certain window, the process moves on, selecting a new window at s distance in time from the previous one. A large ω – *s* overlap guarantees that the evolution of the system is tracked smoothly (i.e. it is convenient to choose *s* << ω).

B.4. Defining the characteristic time scale of the events

Given a certain data set, we do not have any prior knowledge to define the correct timescale Δ*t* at which time series should be aggregated. We do know, however, that activity around civil protests (and in general around any event that involves collective action) are typically far from being stationary or periodic which intuitively points at the fact that there might not exist a single time scale for the whole data set.

Let's consider for now a fixed temporal resolution Δ*t* and bin the Twitter activity such that the first point in the time series will contain any activity that happened between [0, Δ*t*); the second point will contain data from [Δ*t*, 2Δ*t*), and so on. For example, a 30 day data set, sampled at Δ*t* = 60 minutes will render a set of time series of length $L = 30 \times 24 = 720$ points.

As mentioned above, the STE is measured using signals in windows of width ω, spanning from [*t* − ω,*t*). In order to define the optimal Δ*t* we evaluate which temporal resolution provides the best possible information flow among units, i.e. the optimal τ = Δ*t* is the one containing more transfer entropy $s_{\Delta t} = \sum_{x,y} T_{xy}^{\Delta t}$. Operatively, the best timescale is defined as $\{\Delta t : \max_i(s_{\Delta t}^i)\}$. We considered as possible candidates all values Δ*t* from 1 to 120 minutes, with increases of 5 to 15 minutes depending on the data set.

It is important to notice that the maximization of STE at each sliding window may result in different optimal Δt: Our methodological proposal infers from the data the time scale at which events are best described, as illustrated in Figure 2 in the main text.

B.5. Defining the information flow of the events

One may further scrutinize the temporal evolution of the amount of STE each subunit displays. Instead of studying pairwise information flows, in this case we focus on whether a geographical unit is on average driving others or is driven by others at each time step. Within each window of width ω at time t, we calculate the values of the net flow matrix for that window, or normalized directionality index (di), for each geographical unit x at each time step defined as:

$$x_{di} = \begin{cases} \frac{\sum_y T^S_{xy}}{\max_x \sum_y T^S_{xy}} & \text{if } \sum_y T^S_{xy} > 0 \\ \frac{\sum_y T^S_{xy}}{\min_x \sum_y T^S_{xy}} & \text{if } \sum_y T^S_{xy} < 0 \end{cases}$$

Thanks to the normalization $-1 \leq x_{id} \leq 1$ the largest value is associated to the geographical unit exercising the largest driving force to other units. Vice versa, the smallest value is associated to the geographical unit subject to the largest driving forces from other nodes. The result of this measurement --and the corresponding analysis-- can be found in Figure 3 in the main text. Note that each point in the panels of that figure condenses the results obtained for a window integrating information from the *past*, i.e. activity within $[t - \omega, t)$.

B.6. Thresholded versus raw "dominance matrices" $T^\dagger$

Figure 5 in the main text represents θ as measured from thresholded matrices (i.e. matrices in which the lowest Transfer Entropy values have been removed). In this section we discuss fig. S4, where the order parameter θ is represented both for thresholded and raw dominance matrices.

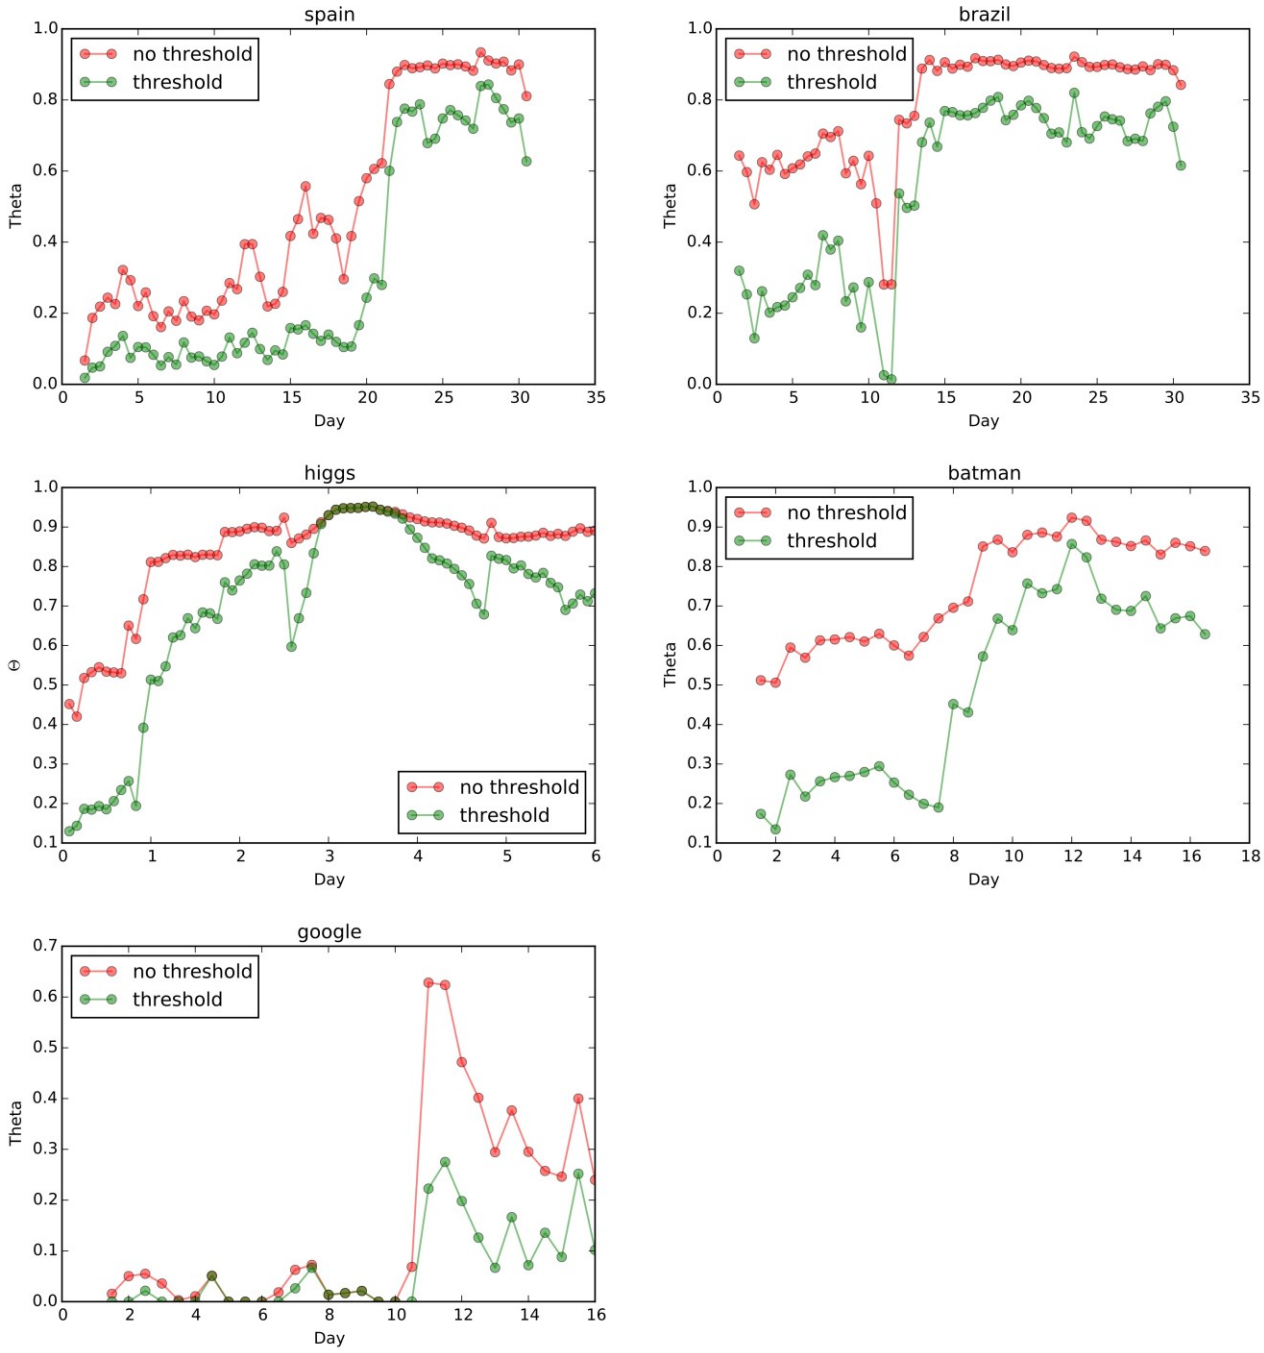

fig. S4. Evolution of the order parameter θ for thresholded (green) and raw (red) $T^{\dagger}$ matrices. Despite a natural increase of the noise in \theta, the interpretation is qualitatively robust.

C. Sensibility analysis of the parametrization

Results in the main text have been obtained with 1-day (i.e., ω = 1) long sliding windows, and embedding dimension m = 5, and using particular geographical units (metro areas in Spain, basins in Brazil, and states in USA). In this section, we present the results of various sensitivity analyses testing how each parameter influences the results.

C.1. The role of the sliding windows width ω

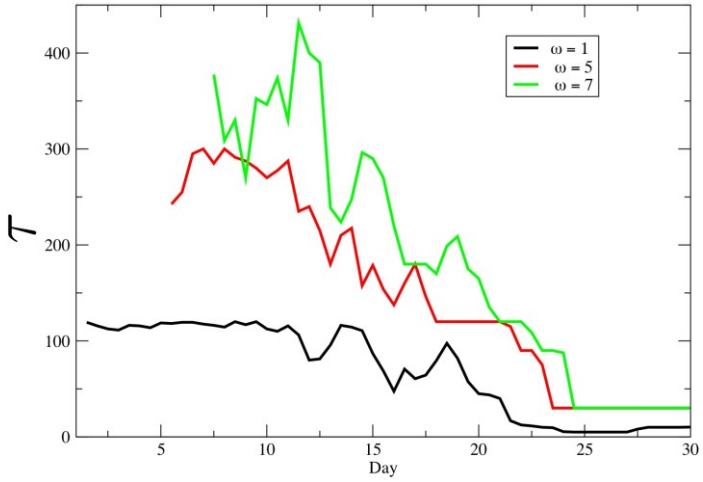

fig. S5. Dependence of τ with he sliding window size ω considering the Spanish 15M protest. Black, red and green lines describe ω = 1,5,7 respectively. Note that, exceptionally, for ω = 5 and ω = 7 days we have considered large Δ*t*, up to 480 minutes (very slow time scales).

Intuitively, if ω is set to a large value the capacity of the method to anticipate events will be reduced, because information emitted long before the present time is affecting the calculations. To check how larger ω's blur the results, we have reproduced the observations in the main text for ω = 5, and ω = 7 given a fixed $m = 5$ considering the 15M data set. Results are offered in fig. S5. It is clear that different window widths behave in the same way as the original one (1 day). Nevertheless, shorter ω yields a more abrupt transition close to the critical event. This observation is more evident studying the behavior of normalized directionality index. In fig. S6 we show this quantity at each time step for the case of the 15M protests in Spain. Similarly to Figure 3 in the main text the size of each bubble is proportional to the logarithm of the activity on twitter. We notice that for ω = 7 and ω = 5 the system shows a change in the driving dynamics clearly before the 15th of May (red strip). Instead, for ω = 1 the transition form a scenario in which the large metropolitan areas are the major driving forces to a more homogenous and delocalized scenario happens during the unfolding of 15th of May.

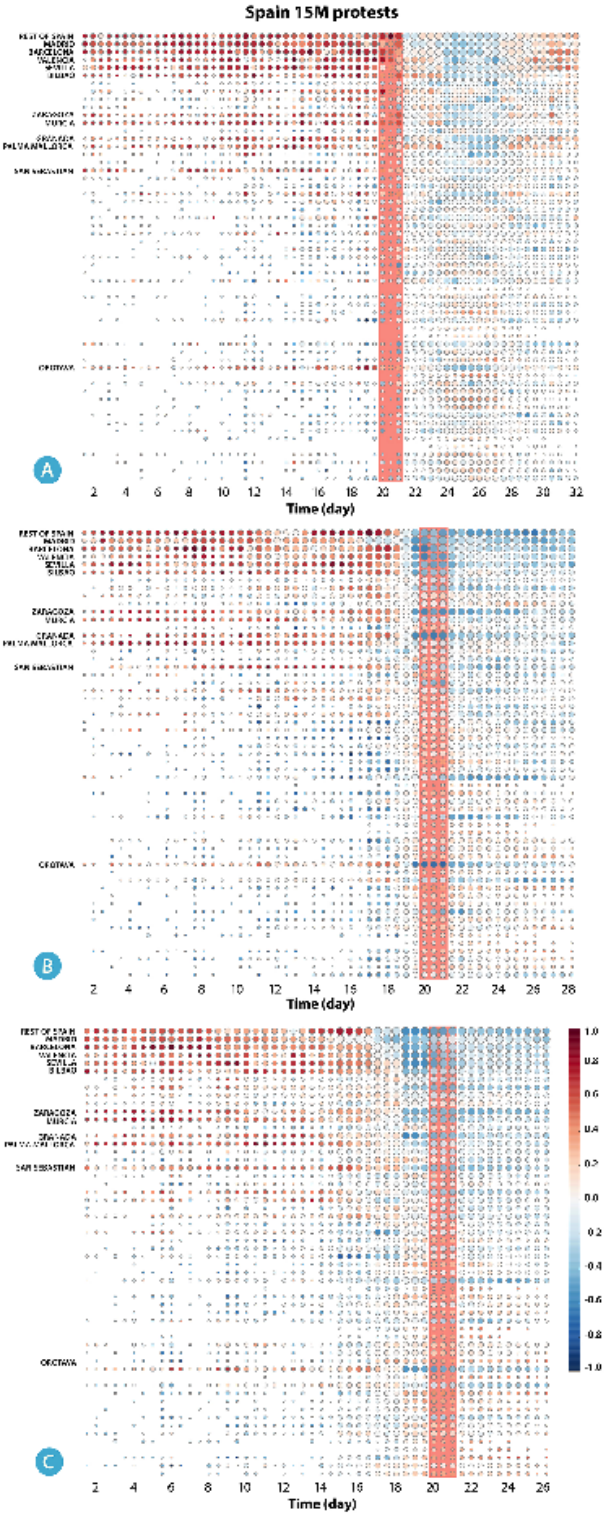

fig. S6. Normalized directionality index for each geographical unit in the 15M data set for different ω. In panel A) we considered ω = 1, in panel B) ω = 5 and in panel C) ω = 7. For all the cases we set $m = 5$. The red strip indicate the 15th of May}

C.2. The role of the embedding dimension m

The embedding dimension m determines how the information in the original time series will be transformed into symbols. The larger m, the larger is the collection of symbols onto which the values are mapped. Since the size of symbols grows like $m!$, it is clear that complex time series demand higher m for a faithful mapping (i.e. one that collects the original complexity). On the other hand, overestimating m adds unnecessary computational costs, because the final result won't change qualitatively. We address the problem of finding the minimal sufficient embedding dimension m, using the approach, called the false nearest neighbor method, proposed by Kennel *et al.* [48].

In practical terms the minimal value of m is found studying the behavior of the nodes encoding the large majority of information. In the case of the 15M protests in Spain this corresponds to Madrid, which was a key spot for the grassroots movements. The minimal value of m is sufficient to disentangle the signal from the dominant node and any other time series in the corresponding data set will need the same or smaller m to be faithfully mapped. Figure S7 reflects these calculations and the strong conclusion is that for any $m \geq 5$ symbolization will have captured the original topology of the real data. Thus, all through the main text, and also in this document, results are reported for $m = 5$, unless indicated otherwise.

In order to further study the effects of m in fig. S8 we plot the behavior of the normalized directionality index for $m = 3, 4, 5$. As it is clear from the plot $m = 5$ (panel C), is able to capture the transition from asymmetric to symmetric scenario in more detail.

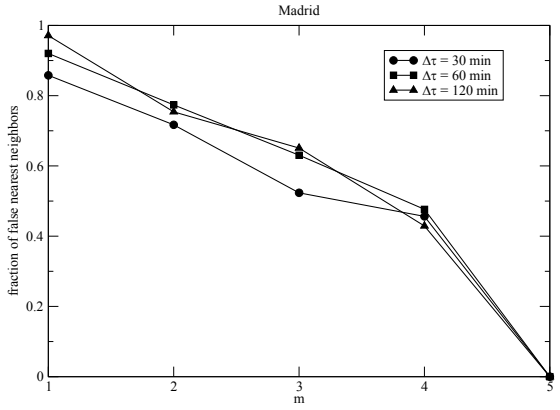

fig. S7. Fraction of false nearest neighbors as a function of m for the Spanish data set and the Madrid time series.

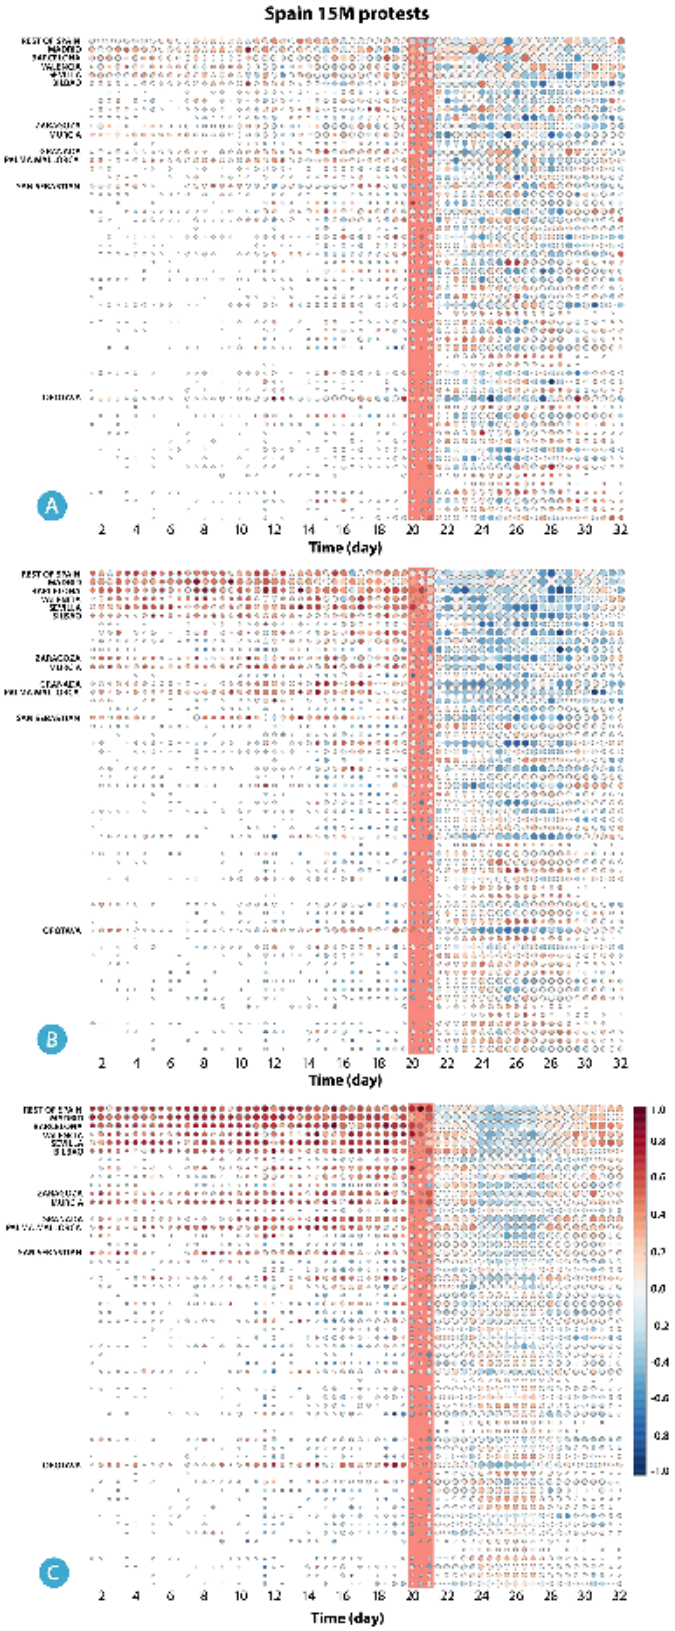

fig. S8. Normalized directionality index for each geographical unit in the 15M data set for different m. In panel A) we considered $m = 3$, in panel B) $m = 4$ and in panel C) $m = 5$. For all the cases we set $\omega = 1$. The red strip indicate the 15th of May.

C.3. Sensibility analysis of the partition

As mentioned above, the Twitter signal could be partitioned in many ways –and we have chosen a geographical approach. In fact, an enormous range of settings are possible: from a simple bipartition of the activity stream to a complete breakdown where a single user is matched to a time series. It is fair then to state up-front that our decision is an arbitrary one, driven by the empirical intuition that geography matters in the emergence of collective action efforts.

Once the geographical partition is adopted, many options are still available: spatial aggregation could be done at the neighborhood, city or county levels (for a finer resolution), or considering a coarser partition. For each data set in the main text we have repeated the analysis for coarser geographical divisions. In the case of Spain's 15M movement, we have moved from the metropolitan areas to the autonomous community level (the Spanish 17 autonomous communities can be regarded as states, i.e. political entities at the regional level [38]). Data from Brazil have been binned in 27 states [37], in contrast with 97 basins in the main text. Finally, US data has been aggregated up to the "divisions" level (9 supra-state areas, as defined by the US Census Bureau [36]).

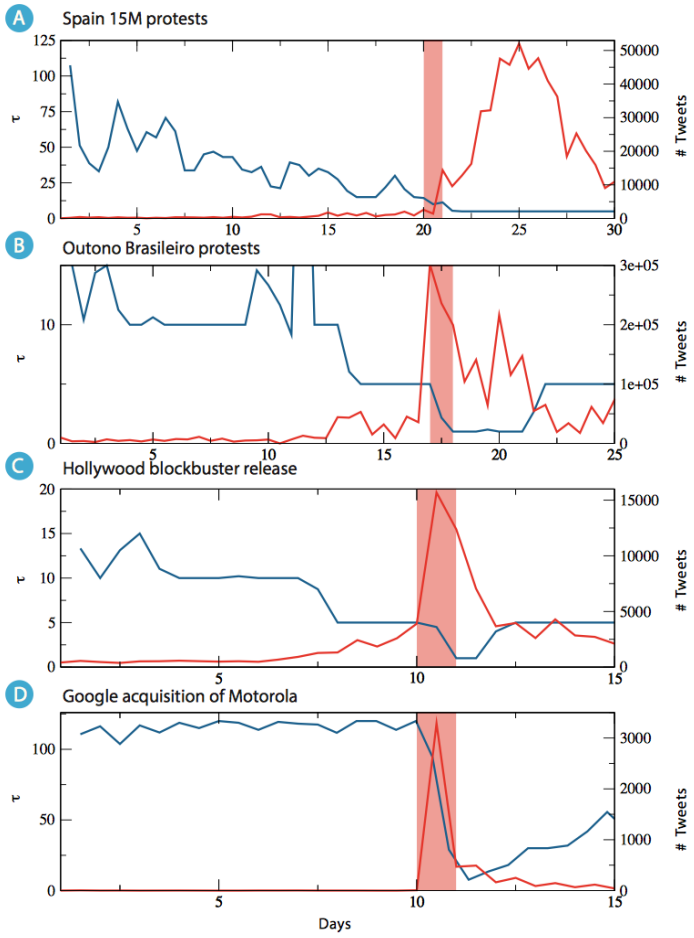

fig. S9. In light blue we plot the characteristic time scale of STE (τ) for the **(A)** 15M protests, **(B)** Outono Brasileiro movements **(C)** release of a Hollywood blockbuster **(D)** the acquisition of Motorola by Google. We considered different geographical aggregations alternative to the ones displayed in

Figure 2 in the main text. The red lines show the activity in Twitter for each data set, and the red strip indicate the day of the main collective event.

Results for these alternative data partitioning can be seen in figs. S9, S10 and S11. Regarding the evolution of the time scales (fig. S9), we observe that the behavior qualitatively resembles the original one in the main text. Time-scales appear to drop at even earlier stages (see fig. S9A); indeed, larger aggregation implies that certain subsystems begin to have a noticeable dynamic role earlier (where previously scattered signal remained insignificant overall). Note however that earlier changes in the time scale come at the cost of a loss in precision, regarding which specific areas are true drivers of the activity (see fig. S11).

A similar result is obtained for the information flow balance in fig. S10, where the occurrence of protests and demonstrations (15M and *Outono Brasileiro*) marks a change in the dominant pattern; the same can be said for the Google-Motorola case.

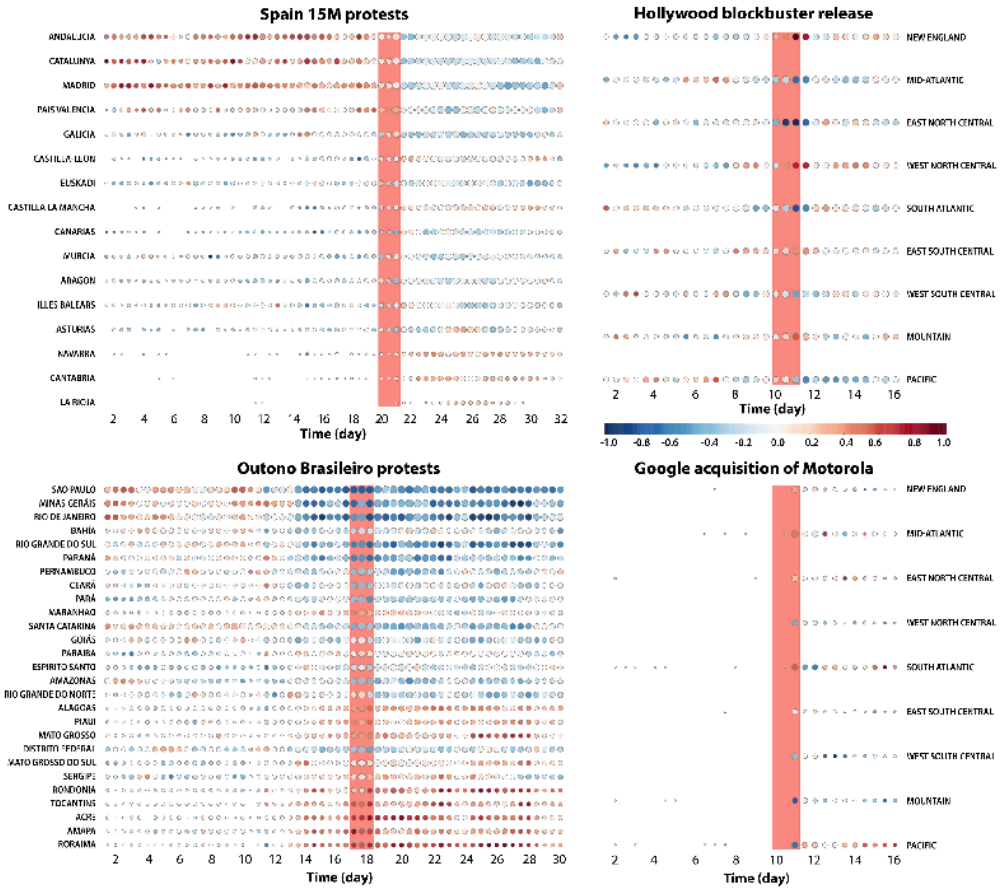

fig. S10. Normalized directionality index for the four data sets aggregated at different geographical levels with respect to those used in the main text. Note that differences with Fig. 3 on the main text are simply due to the new coarsening, and that even in this case the global pattern occurs at similar times and they are just more pronounced.

However, some differences appear in the $\theta - t$ plot (fig. S11). To start with, the Brazilian data set and the Batman event deliver dense T_{xy} matrices, i.e. its sorted counterpart T' has many below-diagonal elements even at early times, indicating that no (or little) transition takes place: the system is decentralized from the very beginning. These differences demand some explanation.

First, it must be highlighted that our tip-over rationale (see Main Text) is valid regardless the apparent contradiction: our claims are concerned with how the values in the T_{xy} matrix are distributed, and as such it is an abstraction of what such matrix represents (be it cities, states or individuals, for that matter). Then, the apparent contradiction simply points at the fact that the lens through which we analyse the events *does* matter. Taking it to the extreme, a bipartition of the data, with two time series accounting for half of the activity each, would easily yield a fully symmetrical T_{xy} matrix; in the opposite situation, a system comprising each user individually would render an (almost) empty matrix, given the fact that most people is not showing activity most of the time.

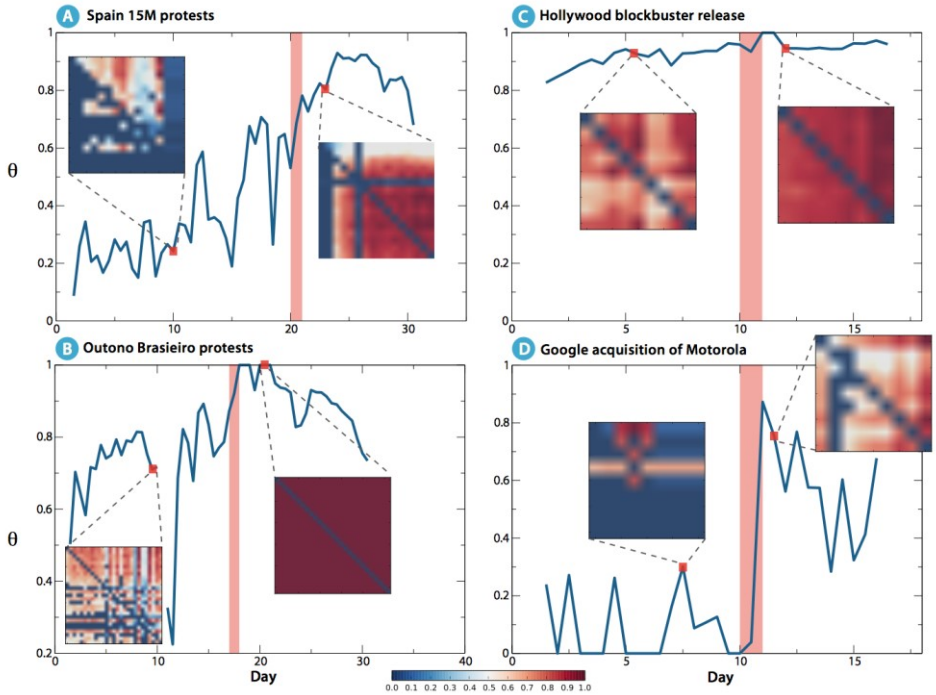

fig. S11. Behavior of θ as a function of time for different geographical aggregations. For each data set two matrices T' are plotted considering a time before and one after the main event (signaled with a red vertical red bar). A blurred transition can still be observed for events A and B. Note a point missing in the Brazilian data set due to a data blackout between days 10 and 11.

All in all, these results suggest that our proposal opens up exciting research questions: for instance, at which level of resolution should the system be observed to extract an optimal analysis out of it? We must keep in mind that other relevant events in Twitter do not have a geographical component; groups may be defined by religious beliefs, age, gender issues, or other demographic attributes. It remains beyond the scope of this work to determine how to obtain optimal partitions that will render the correct conclusions. For the time being, we rely on commonsensical, predefined –rather than optimally detected– entities (geographical, in this case) to make a case of our methods and rationale.

D. Validation of results (I): Time series randomization

In this section we validate our framework studying its performance on data surrogates, i.e. statistical ensembles of randomized data. In other words, we apply our approach to a set of data that by construction do not contain the temporal correlations we find in real data sets. This step is crucial to prove that our observations capture genuine features of real collective events.

D.1. Statistical randomized surrogates of original data

In order to validate our results, we need to make sure that our analysis and conclusions are not mere artifacts which would arise in any case. To provide a reasonable baseline, we need to build randomized counterparts of the data and then analyze it just as we did for the actual case. We consider two methods to obtain data random surrogates: amplitude adjusted Fourier transform surrogates and constrained randomization surrogates, with an extensive use of the TISEAN software [50].

We present the results for the randomization of two data sets, with qualitatively similar insights for the other data sets.

D.2. Amplitude Adjusted Fourier Transform (AAFT) surrogates

A first, robust step to provide a suitable null model is to generate randomized data sets which ensure that certain features of the original data will be preserved. In particular, we generate AAFT surrogates as proposed in [51], who established an algorithm to provide surrogate data sets containing random numbers with a given sample power spectrum and a given distribution of values.

Under these constraints, we obtained 50 randomized versions of the 15M data, which were then analyzed in the same way as the original data (see Main Text and previous sections). The averaged results from such analysis are offered in fig. S12. Clearly, the original patterns are completely blurred and just a single characteristic time scale can be observed. Furthermore, our approach do not capture any change in the characteristic time scale as correlations and driving between different units have been artificially eliminated in the data.

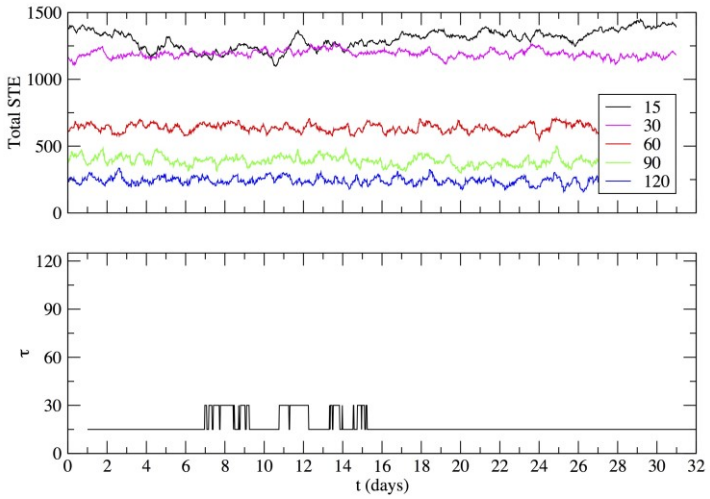

fig. S12. Average total amount of STE for some Δt (top panel) and time scale profile τ (bottom panel) for 15M amplitude adjusted Fourier transform surrogates (50 randomizations).

Mirroring our analysis of real data, we intend to see whether some trace of the original transition is kept for this newly obtained random version of the data. To do so, we also exploit 50 randomizations of the original data set, for which we can extract average surrogate snapshots (i.e. the state of the system at a given *t* day). Just as in Figure 5 of the main text, fig. S13 shows two sorted (ranked) $T^{\dagger}$ matrices, corresponding to two different moments (before and after the main event) for the statistical randomized surrogates of original data. It is interesting to notice that as any localized abrupt change in the time-scale is washed out (fig. S12), also any sort of systemic transition is missing.

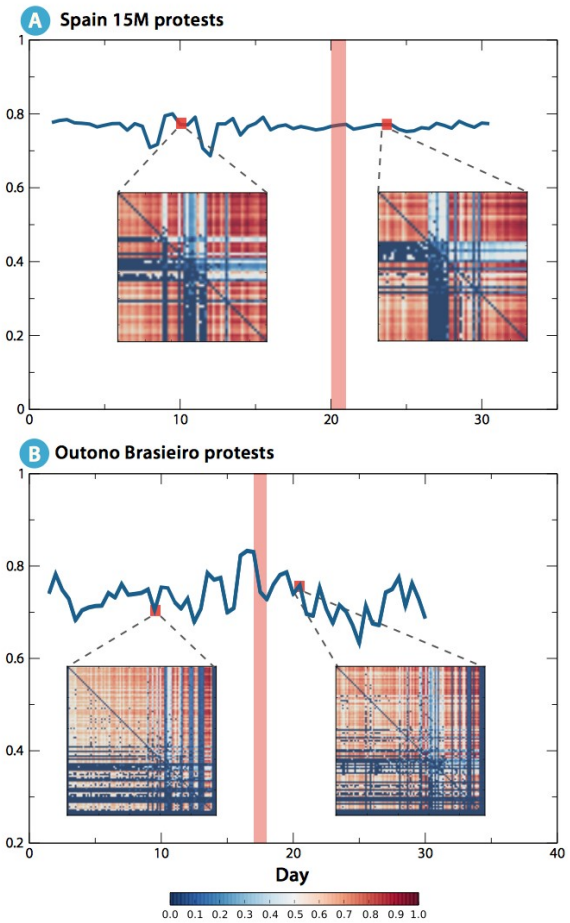

fig. S13. Behavior of θ as a function of time for the statistical randomized surrogates of the 15M **(A)** and Brazilian data set **(B)**. For each data set two matrices $T^{\dagger}$ are plotted considering a time before and one after the main event (signaled with a red vertical red bar). As expected in this context, no significant transition is observed in neither cases.

D.3. Constrained randomization surrogates

Beyond a randomization scheme that guarantees given power spectrum and distribution of values, one might want to generate surrogates which are further constrained. This can be achieved if we demand randomized data sets to preserve as well a given non-periodic autocorrelation function (ACF). To this end, Schreiber [51] developed a method of constrained randomization of time series data which seeks to meet the given constraints through minimization of a cost function, among all possible permutations, by the method of simulated annealing.

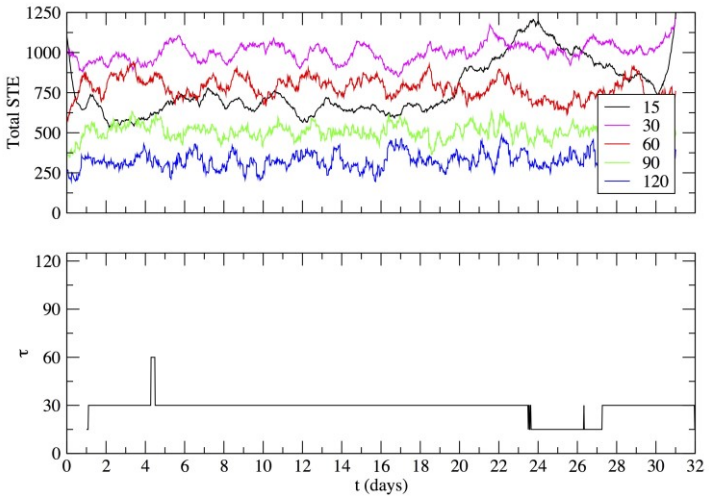

fig. S14. Average total amount of STE for some Δt (top panel) and time scale profile τ (bottom panel) for 15M constrained surrogates (20 randomizations).

In fig. S14 the averaged results for the analysis of the constrained surrogate data can be checked. Twenty randomizations for the Spanish data set were obtained. Even with the additional constraints (if compared with AAFT randomizations, see previous section), hardly any resemblance with the original patterns can be observed. It must be noted that constrained randomizations are time and CPU-consuming, due to the additional restrictions regarding ACF.

E. Validation of results (II): Unfiltered Twitter stream

To further test the contributions of the work –a methodological framework to reveal the evolving time scales of collective events, and its inherent order-disorder transitions– we entertained the idea of analyzing "less successful events". However, such an idea is difficult to implement –if events are less successful, most likely it is not possible to know which are the appropriate words to filter the Twitter stream. Therefore, we propose the opposite: we have collected an *unfiltered* sample of geocoded tweets from all around the world, and analyzed its development for some time.

Some facts:

- Only 2 days have been analyzed (enough to grasp any pattern, if it exists, given that we are not monitoring any specific event). This roughly corresponds to ~3M tweets with precise geolocation.

- We have monitored the stream with window width $\omega = 6$ hours (note this is a smaller window than for our previous analyses, where $\omega = 24$h).

- Each 6 hour window "advances" 30 minutes to analyze the next snapshot (this has not changed from our previous analyses)

- The embedding dimension has been kept at $m = 5$.

- Original Twitter data contains tweets from 230 countries. Tweets from US are most frequent, accounting for ⅓ of the total (~1M); Niue is least frequent (1 tweet). To keep computational costs at a reasonable level, we have arbitrarily thresholded the list of countries. We have set that threshold at 5000 tweets (in the whole period), retaining only 42 countries.

- The set of Δt values that have been used is: 1, 5, 10, 15, 30, 45, 60, 120, 180, 300 and 600 seconds (in brief: a good sample between 1 sec and 10 min). Note that in our case studies we were in the order of minutes, not seconds (in unfiltered scenarios, we foresee the time scales to be very fast).

Results:

The resulting STE-optimal time scales (τ) analysis yields an almost constant value (45 sec) with a few ticks in which it is 30 sec, see fig. S15 below. This is indeed expected. Additionally, this finding strengthens the idea that large volumes are sufficient to derive fast time scales. The fact that we find fast time scales in the *absence* of large volumes enhances the value of the proposed tools.

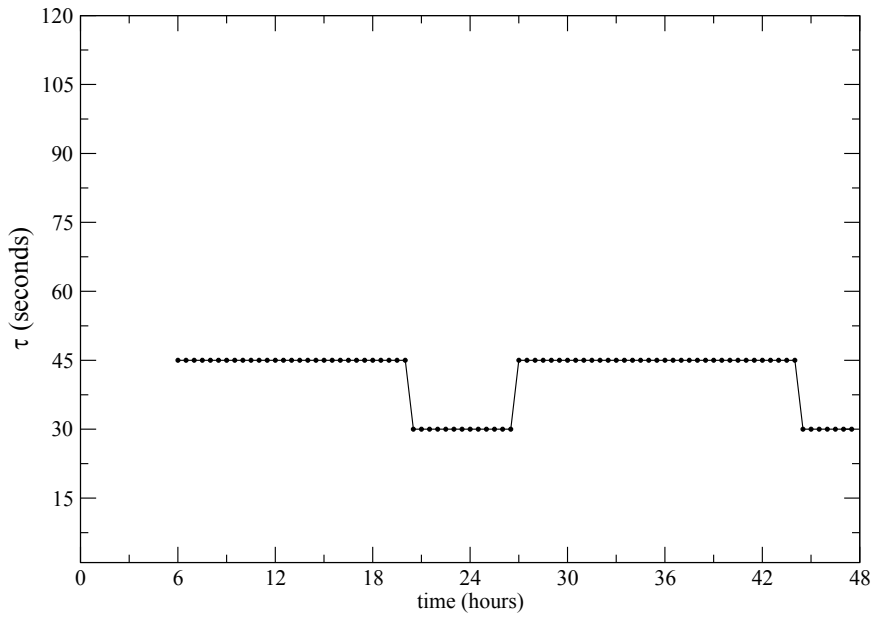

fig. S15. Evolution of τ as a function of time. In an unfiltered scenario, there are no distinguishable time scales to characterize the system's development. For the sake of clarity we represent here up to 120 sec on the y-axis, despite calculations were made up to 600 seconds.

In fig. S16 below we see two thresholded $T^{\dagger}$ matrices corresponding to different moments in the unfiltered data set: the very first one (from 0 to 6h), and one from the final part of the studied period (from 36h to 42h). These are good representatives of any of the snapshots, in which clearly no pattern is to be observed (i.e. even the row/column re-ordering and the thresholding do not allow to infer clear driven/driving roles). It is worth remarking again that this figure corresponds to a thresholded matrix, i.e. the version that should ease visualization if any patterns existed.

fig. S16. Thresholded $T^{\dagger}$ matrices corresponding to different moments in the Twitter unfiltered data set. As expected, no country-to-country driving patterns emerge.

Additionally, to have a complete picture, we include two figures (see below) with the actual time series (volume). Each set of plots (left, right) has 4 panels: global volume, volume for US (top contributor), Kuwait (the country in the middle, in terms of volume), and Nigeria (last included, in terms of volume).

The set of plots on the left represent raw volume at a resolution of 600 seconds (less noisy), the one on the right at the dominating time scale (45 sec).

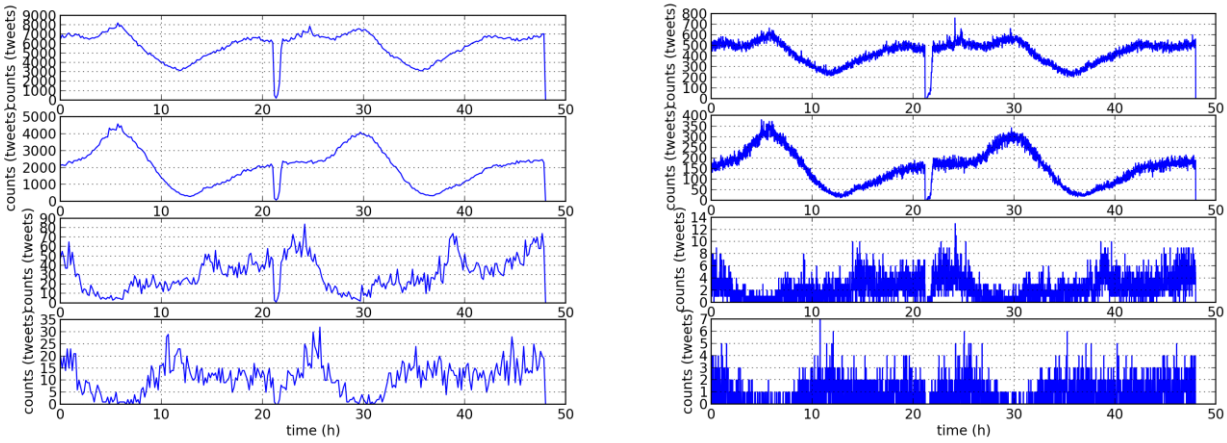

fig. S17. Raw time series for Twitter unfiltered stream for Δt = 600 s and Δt = 45 s (left and right, respectively). In both cases, global (top), US (second), Kuwait (third) and Nigeria (bottom) time series are represented.

Some interesting insights can be drawn from the raw time series:

- In general, we observe circadian rhythms for these countries (2nd to 4th panel); the global panel (top) also shows such rhythm --as it is dominated by US.

- There is a data gap at t ~21h, spanning approximately one hour. This is perhaps due to a storage problem in the Decahose system where the data were collected from. Because we always consider $\omega = 6$ hours, overall results are not affected.

- Because we are dealing with an unfiltered data set, the circadian rhythms are naturally shifted (different time zones). Note that Kuwait and Nigeria (3rd and 4th panels) seem to be advanced with respect with US (i.e. the trend in Kuwait at $t = 0$ is the same as US for $t = \sim 6$). It is important to highlight that this shift can imply that STE detects spurious driving (which would translate into hot colors in a plot similar to those in Figure 3 of the main text). However, this doesn't affect the timescales, which are low and relatively constant whatsoever. And, while an information flow evolution plot (similar to Figure 3 in the main text) might wrongly represent spurious driving, the "dominance matrices" $T^{\dagger}$ (fig. S16) are telling us precisely that no real driving is actually happening. That's precisely why the order parameter θ is absolutely relevant.

Finally, we would like to stress that the window width choice for this counterexample (ω = 6h) represents a harder test for our framework, because a smaller time window should detect shorter events (say, the rise and fall of some news within hours).

F. Validation of results (III): Synthetic time series generation

This work is focused on the capabilities of STE as a viable tool to analyze empirical data. However, it is under a synthetic controllable setting that the capabilities and findings of our approach can be made explicit –that is, the existence of transitions in the time scales (slow-to-fast) and directional couplings (asymmetric-to-distributed) in systems that gather around some sort of collective action.

Here, we offer two synthetic examples to better show the relationship between the optimal time scale (τ), the overall activity (or volume) and the dynamic interaction. These, as entangled as they may seem in our empirical examples, are here clearly separated.

Part 1) Minimalist example: Disentangling volume and time scales (Δt).

The question: can we imagine two constant-volume scenarios that maximize STE at different Δt? The answer is yes:

Take a time series Y = constant (e.g. all 0s, or all 1s, it won't make a difference)
Now imagine two time series X_i that have the same activity volume:

$$X_1 = \{1\ 0\ 1\ 0\ 1\ 0\ 1\ 0\}$$
$$X_2 = \{1\ 1\ 0\ 0\ 1\ 1\ 0\ 0\}$$

(each tick is a generic time unit, and we test $\Delta t = 1$ and $\Delta t = 2$, with m = 2)

Note that the volume is constant system-wide and even subsystem-wise (i.e. overall activity is constant, and is constant also if you check individual time series).

Results:

<u>System X_1,Y</u>

$$\Delta t = 1 \qquad T_{x1,y} = \begin{pmatrix} 0 & .014 \\ 0 & 0 \end{pmatrix}$$

$$\Delta t = 2 \qquad T_{x1,y} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$$

We infer from this that the optimal time-scale in terms of STE is $\tau = 1$

System X_2,Y

$$\Delta t = 1 \qquad T_{x2,y} = \begin{pmatrix} 0 & .007 \\ 0 & 0 \end{pmatrix}$$

$$\Delta t = 2 \qquad T_{x2,y} = \begin{pmatrix} 0 & .08 \\ 0 & 0 \end{pmatrix}$$

We infer from this that the optimal time-scale in terms of STE is $\tau = 2$

Therefore it is proved that *it is the dynamical interplay between the system's units that modifies the optimal time scale*. While the opposite is not always true: changes in volume do not *necessarily* mean changes in the dynamical interaction: just think of *Y* as an always-increasing time-series (instead of constant), and the result would be exactly the same (because *Y* would be translated into a single symbol, i.e. a constant symbolized series, rendering again the reported results).

Under this extremely simple example there is an implicit idea: one could come up with a synthetic model –which would allow for extensive testing in unlimited scenarios– if there were some efficient way to generate time series with some prescribed permutation entropy (which is beyond the scope of the work, but see ref. [22]).

Part 2) Nonlinear Lorentz oscillators: Time scales, volume, and dynamical coupling

Another way to approach the entanglement between time scales, volume, and dynamic interaction is to revisit the original Symbolic Transfer of Entropy paper by Staniek & Lehnertz [21]. In their paper, the authors propose a couple of synthetic benchmarks on which they demonstrate STE's virtues. It seems then natural to get inspiration from them here too.

Here, we monitor two Lorentz oscillators with varying levels of coupling, and see if changes in the time scale are observed. For our purposes, there are two stages we are interested in (but the possibilities are virtually infinite): (a) oscillators *A* and *B* are completely independent from each other (uncoupled, $\beta = 0$); (b) we increase the coupling strength such that oscillator *A* drives *B* (asymmetric coupling, $\beta > 0$. The results can be seen in fig. S18.

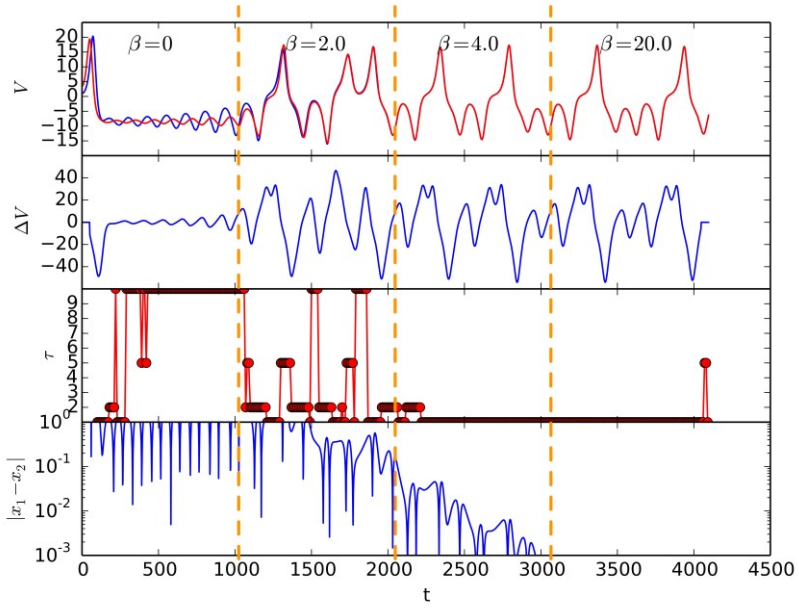

fig. S18. Evolution of two nonlinear systems under four changing scenarios: from dynamic independence (β = 0) to strong asymmetric coupling (β = 20.0).

This figure represents the evolution of the x-dimension of both oscillators for 4096 integration steps. As indicated by vertical lines, the unidirectional coupling strength β from oscillator *A* to *B* has been progressively increased, each 1024 steps. STE has been calculated for the whole system, with $w = 100$, $m = 5$ and different resolution levels, $\Delta t = \{1, 2, 5, 10\}$.

In the top panel we can see the evolution of both oscillators (blue: *A*; red: *B*). The second panel shows the variation in volume from window to window ($\Delta V_t = V_{w(t)} - V_{w(t-1)}$). The third panel shows the timescale at which STE is optimal (i.e., τ). Finally, the fourth plot shows the absolute difference between the values from oscillators A and B, as a (rough) proxy for their level of synchronization (note the log scale).

Some lessons from the plot can readily be extracted. First, the time scale is very sensitive to the coupling of the subsystems. When *A* and *B* are uncoupled, the optimal time scale is the slowest (τ = 10); the system immediately reacts to the introduction of slight coupling strength β, the time scale drops to 2 – then behaves erratically to some extent, as the system is not yet strongly coupled nor fully synchronized. Note that the sudden decrease in τ appears to happen slightly after the coupling strength is increased: this is an artifact of our representation, as we always report results for the *past* sliding window. That is, each red dot in the plot represents the corresponding τ for the *previous w* points of the time series.
For the stages with large β both systems become fully synchronized and the time scale remains low. *It is the dynamic relationship between both oscillators that determines* τ.

It is interesting to see the evolution of Δ*V* and τ (2nd and 3rd panels). The "variation in volume" Δ*V* is mostly constant during the first stage; after it, volume is rapidly changing (large Δ*V*). Out of visual inspection, it is possible to observe decreases in τ when volume is growing (Δ*V* > 0); and yet, it is

possible to observe also the opposite, i.e. τ *can decrease with volume decreases* ($\Delta V < 0$). The straightforward relationship (increased activity implies short timescales, and vice versa) is –now via dynamical models– refuted.

On a noisy τ: During the second stage of the dynamics (unidirectional weak coupling, $\beta = 2.0$) it is possible to observe some fluctuations of the optimal time scale τ. These can be explained from the chaotic trajectories of the units, which, under the mentioned weak regime, sometimes are close to synchronization (and sometimes not: note also the fluctuations in the bottom panel). STE (or any other measure, for that matter) is incapable of disentangling these different scenarios and reflects, at this very local level (keep in mind that $w = 100$ only). On the positive side, this brings some light to the accuracy and sensitivity of the measure; on the downside, it provides at times a noisy result –which nonetheless is, to our understanding, quite robust if observed for a longer time span. That is, if we observe the overall differences in τ for the periods $t < 1024$ and $1024 < t < 2048$, it becomes clear that the time scale has undergone a transition.

Having said that, we finally wish to highlight that these local fluctuations can be smoothened if a longer w is considered: see fig. S5. This comes at the cost of a less sensitive measure –the time scale decrease is observed at later stages.