

The spontaneous emergence of conventions: An experimental study of cultural evolution

Damon Centola^{a,b} and Andrea Baronchelli^{c,1}

^aAnnenberg School of Communication, University of Pennsylvania, Philadelphia, PA 19106; ^bCenter for Advanced Study in the Behavioral Sciences, Stanford University, Palo Alto, CA 94305; and ^cDepartment of Mathematics, City University London, London EC1V 0HB, United Kingdom

Edited by Giorgio Parisi, University of Rome, Rome, Italy, and approved December 23, 2014 (received for review September 30, 2014)

How do shared conventions emerge in complex decentralized social systems? This question engages fields as diverse as linguistics, sociology, and cognitive science. Previous empirical attempts to solve this puzzle all presuppose that formal or informal institutions, such as incentives for global agreement, coordinated leadership, or aggregated information about the population, are needed to facilitate a solution. Evolutionary theories of social conventions, by contrast, hypothesize that such institutions are not necessary in order for social conventions to form. However, empirical tests of this hypothesis have been hindered by the difficulties of evaluating the real-time creation of new collective behaviors in large decentralized populations. Here, we present experimental results—replicated at several scales—that demonstrate the spontaneous creation of universally adopted social conventions and show how simple changes in a population’s network structure can direct the dynamics of norm formation, driving human populations with no ambition for large scale coordination to rapidly evolve shared social conventions.

social conventions | spontaneous emergence | complex systems | empirical testing | network science

Social conventions are the foundation for social and economic life (1–7). However, it remains a central question in the social, behavioral, and cognitive sciences to understand how these patterns of collective behavior can emerge from seemingly arbitrary initial conditions (2–4, 8, 9). Large populations frequently manage to coordinate on shared conventions despite a continuously evolving stream of alternatives to choose from and no a priori differences in the expected value of the options (1, 3, 4, 10). For instance, populations are able to produce linguistic conventions on accepted names for children and pets (11), on common names for colors (12), and on popular terms for novel cultural artifacts, such as referring to junk email as “SPAM” (13, 14). Similarly, economic conventions, such as bartering systems (2), beliefs about fairness (3), and consensus regarding the exchangeability of goods and services (15), emerge with clear and widespread agreement within economic communities yet vary broadly across them (3, 16).

Prominent theories of social conventions suggest that institutional mechanisms—such as centralized authority (14), incentives for collective agreement (15), social leadership (16), or aggregated information (17)—can explain global coordination. However, these theories do not explain whether, or how, it is possible for conventions to emerge when social institutions are not already in place to guide the process. A compelling alternative approach comes from theories of social evolution (2, 18–20). Social evolutionary theories maintain that networks of locally interacting individuals can spontaneously self-organize to produce global coordination (21, 22). Although there is widespread interest in this approach to social norms (6, 7, 14, 18, 23–26), the complexity of the social process has prevented systematic empirical insight into the thesis that these local dynamics are sufficient to explain universally adopted conventions (27, 28).

Several difficulties have limited prior empirical research in this area. The most notable of these limitations is scale. Although compelling experiments have successfully shown the creation of

new social conventions in dyadic and small group interactions (29–31), the results in small group settings can be qualitatively different from the dynamics in larger groups (*Model*), indicating that small group experiments are insufficient for demonstrating whether or how new conventions endogenously form in larger populations (32, 33). Important progress on this issue has been made using network-based laboratory experiments on larger groups (15, 24). However, this research has been restricted to studying coordination among players presented with two or three options with known payoffs. Natural convention formation, by contrast, is significantly complicated by the capacity of individuals to continuously innovate, which endogenously expands the “ecology” of alternatives under evaluation (23, 29, 31). Moreover, prior experimental studies have typically assumed the existence of either an explicit reward for universal coordination (15) or a mechanism that aggregates and reports the collective state of the population (17, 24), which has made it impossible to evaluate the hypothesis that global coordination is the result of purely local incentives.

More recently, data science approaches to studying norms have addressed many of these issues by analyzing behavior change in large online networks (34). However, these observational studies are limited by familiar problems of identification that arise from the inability to eliminate the confounding influences of institutional mechanisms. As a result, previous empirical research has been unable to identify the collective dynamics through which social conventions can spontaneously emerge (8, 34–36).

We addressed these issues by adopting a web-based experimental approach. We studied the effects of social network structure on the spontaneous evolution of social conventions in populations without any resources to facilitate global coordination

Significance

Social conventions shape every aspect of our lives, from how we greet each other to the languages we speak. Yet their origins have been a topic of theoretical speculation since the time of Aristotle. Most approaches assume that institutions are necessary to organize large populations, but the simplest explanation is that universally accepted conventions are the unintended consequence of individuals’ efforts to coordinate locally with one another. Although this hypothesis is compelling, it lacks conclusive empirical support. Here, we present results from controlled experiments demonstrating that changes in network connectivity can cause global social conventions to spontaneously emerge from local interactions, even though people have no knowledge about the population, or that they are coordinating at a global scale.

Author contributions: D.C. and A.B. designed research, performed research, analyzed data, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. Email: a.baronchelli.work@gmail.com.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1418838112/-DCSupplemental.

(9, 37). Participants in our study were rewarded for coordinating locally, however they had neither incentives nor information for achieving large scale agreement. Further, to eliminate any preexisting bias in the evolutionary process, we studied the emergence of arbitrary linguistic conventions, in which none of the options had any a priori value or advantage over the others (3, 23). In particular, we considered the prototypical problem of whether purely local interactions can trigger the emergence of a universal naming convention (38, 39).

Theoretical Model

The approach used here builds on the general model of linguistic conventions proposed by Wittgenstein (39), in which repeated interaction produces collective agreement among a pair of players. Theoretical extensions of this approach have argued that myopic players interacting in social networks can unintentionally create percolating cascades of coordinated behavior (6, 10, 23, 25, 27, 40, 41). Theoretical predictions for our study are based on a derived “language game” model of convention formation (27), in which agents attempting to coordinate in pairwise interactions accrue a memory of past plays, which they use to “guess” the words that will be used by their subsequent partners (*Model*). Consistent with a broad range of formal approaches (5, 33, 42–44), this model predicts that the connectivity of the actors’ social networks can influence the collective dynamics of convention formation, ranging from the emergence of competing regional norms that inhibit global coordination (45) to the rapid growth of universally shared social conventions (27) (*Model*).

We evaluated these predictions by studying convention formation in three representative network configurations: (i) spatially embedded social topologies (i.e., one-dimensional lattices with degree 4) (45, 46, 47), (ii) randomly connected topologies (i.e., random graphs with constant degree 4) (42, 48), and (iii) homogeneously mixing populations (3, 27). Formal results show that alternative network configurations (48) fall within the range of dynamical behavior exhibited by the three topologies used here (42).

Experimental Design

Each live game, or experimental “trial,” consisted of a set of participants, a specific social network structure, and a prespecified number of rounds to play. When participants arrived to play the game, they were randomly assigned to positions within a social network. In a given round of the game, two network “neighbors” were chosen at random to play with one another. Both players simultaneously assigned names to a pictured object (i.e., a human face), blindly attempting to coordinate in the real-time exchange of naming choices (Fig. S2). If the players coordinated on a name, they were rewarded with a successful payment; if they failed, they were penalized (*Materials and Methods*). After a single round, the participants could see only the choices that they and their partner had made, and their cumulative pay was updated accordingly. They were then randomly assigned to play with a new neighbor in their social network, and a new round would begin. The object that participants were trying to name was the same for the entire duration of the game, and for all members of the game. Participants in the study did not have any information about the size of the population that was attempting to coordinate nor about the number of neighbors to whom they were connected.

Results

Fig. 1 ($n = 24$) shows that the dynamics of emergent social conventions depend decisively upon the structure of the social network. In spatial networks, populations enjoyed rapid local coordination, in some cases achieving a 50% success rate (i.e., average likelihood of matching words with a partner) as early as round 4. However, this initial success rate quickly decelerated. After 25 rounds of play, average success rates failed to reach above 75%. Throughout all of the spatial network trials, the dominant local conventions (i.e., the most popular word choices) were never used by more than 45% of the population. As shown in Fig. 2, behavior in the spatial networks evolved through a process of local coarsening (45), in which emergent regions of coordinated behavior competed with bordering local conventions (45, 49, 50). In each of the trials, these dynamics inhibited the spontaneous emergence of global coordination by

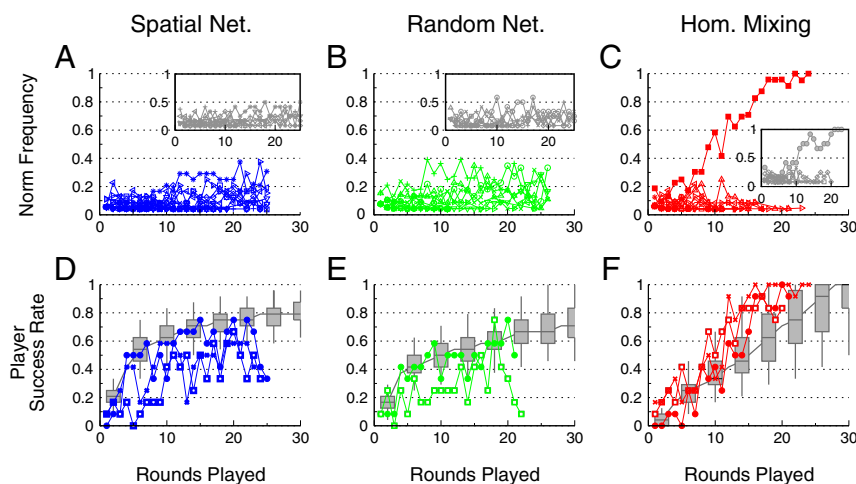


Fig. 1. The evolution of social conventions across network topologies ($n = 24$). Panels indicate spatially embedded networks (A and D, blue), random networks (B and E, green), and homogeneously mixing populations (C and F, red), for eight independent trials of the study ($n = 24$). Panels A–C show the evolving ecology of norms for representative trials from each condition. Insets show representative model simulations. The corresponding time series (D–F) show the average level of successful matching among individual players. Model results are shown in gray (95% confidence intervals over 10,000 realizations). In spatially embedded networks (A and D), players achieved moderate success with local conventions, creating regional competition and preventing a single convention from emerging across the population. Similarly, in random networks (B and E), moderately successful local coordination produced groups in the networks, but no global consensus. By contrast, in homogeneously mixing populations (C and F), initial local failures resulted in rapid population-level learning and global coordination on a single convention.

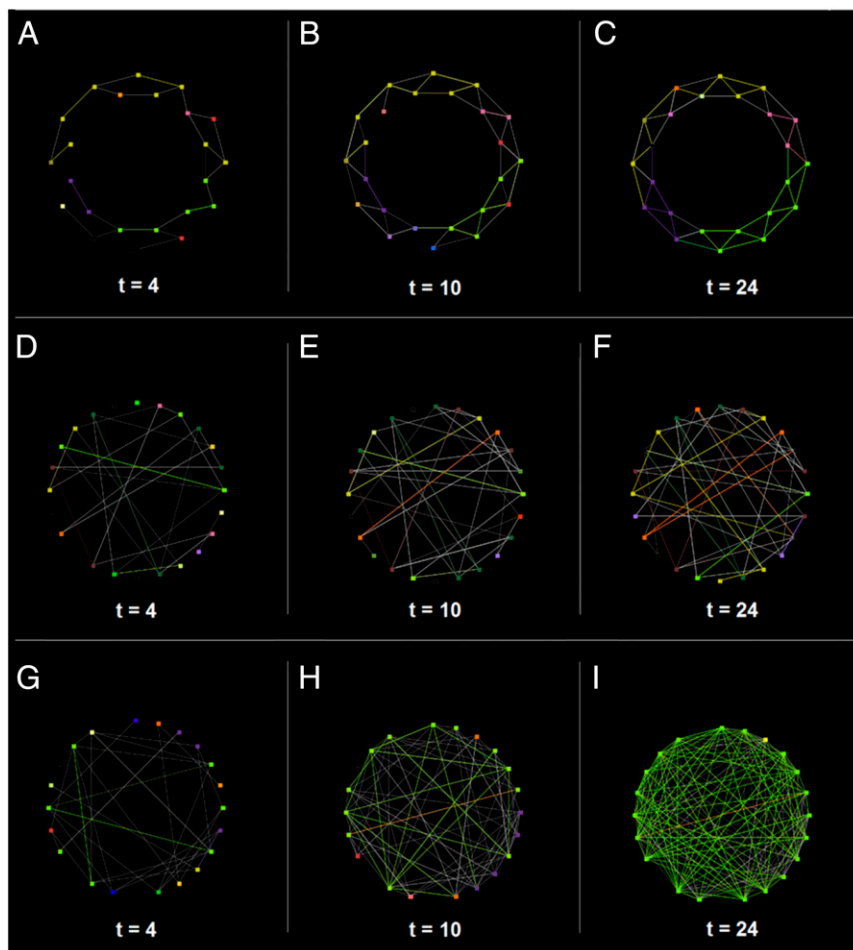


Fig. 2. Temporal dynamics of convention formation in representative experimental trials. Panels correspond to a spatially embedded network (A–C), a random network (D–F), and a homogeneously mixing population (G–I) ($n = 24$). Each color corresponds to a unique name used in the trial. Node color refers to the name that was most recently used by that actor, and edge color indicates the name that the two players matched on in their most recent interaction. A white edge indicates that the two players failed to match in their most recent interaction. In the spatial network (A–C), local interactions produce clusters of coherent coordination around a shared convention, with contested border regions. A similar dynamic unfolds in the random network (D–F), where repeated interaction leads to local coordination. In the homogeneously mixing population (G–I), a single name assumes dominance, becoming the global convention.

creating entrenched competition between endogenously formed groups (51).

Similar results were found in random networks, in which local groups of coordinated individuals emerged and competed for dominance (Fig. 1). Group formation in random networks was driven by repeated interactions, which created metastable boundaries between groups of neighbors despite the absence of local clustering (Fig. 2). After 25 rounds of play in randomized topologies, local groups still persisted and coordination rates never increased above 75%. In all random network trials, global social conventions never emerged. Moreover, the sizes of the dominant social conventions (i.e., the fraction of the population using the most popular word choice) were equivalent across all trials of the spatial and random networks, averaging 33% of the population. On time scales observable within our study, the dynamics of social coordination in both the spatial and random network trials were driven by local group competition, which impeded the emergence of global conventions.

Homogeneously mixing populations exhibited significantly different dynamics than those observed in the other two topologies. Initially, success rates were lower because actors did not have repeated interactions with their partners, which prevented “neighborhoods” of entrenched behavior from forming. However, local failure accelerated global coordination. In all trials

with homogeneously mixing populations, success rates increased to 100% well before the end of the study. Fig. 1 shows that this rapid growth in individual success corresponds to the spontaneous emergence of a global social convention. In all trials, an emergent convention grew quickly, reaching over 60% of the population by round 12, and achieving universal adoption between rounds 20 and 22. On average, by round 22, players who had never interacted with one another were all using the same convention and were able to consistently coordinate with new partners.

The speed of self-organized conventions in these networks raises the question of whether these coordination dynamics scale up as population sizes increase. There are good reasons for skepticism. As system size increases, so does the expected number of competing alternatives circulating in the population; at the same time, because interactions are limited by the number of rounds in the game, increasing system size reduces the fraction of the population with which any given individual can interact. These considerations suggest that global coordination may be much more difficult in larger populations.

We tested this conjecture by doubling the size of the population and replicating our study. These larger trials ($n = 48$) permit a more detailed view of the evolving competitive landscape that constitutes the “ecology” of social conventions. Fig. 3 shows the changing distribution of popularity among the competing alternatives in all

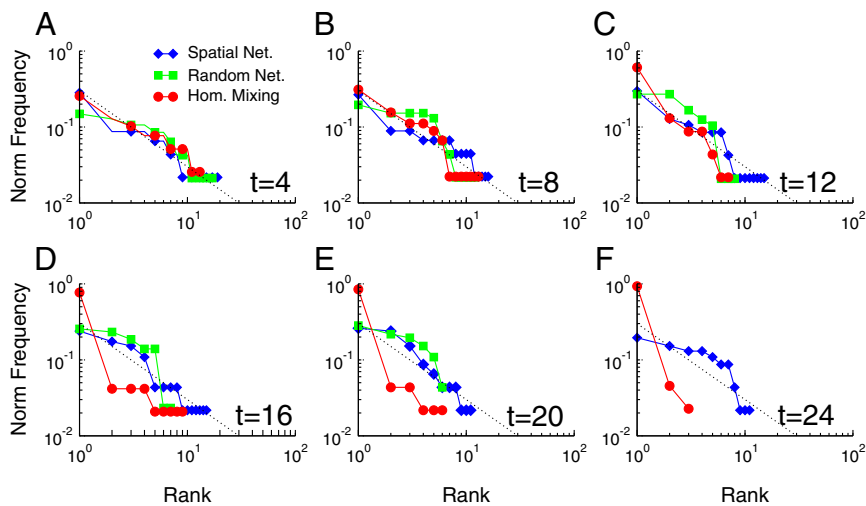


Fig. 3. The shifting distributions of conventions in diverse network topologies. Frequency-rank plots for spatially embedded networks (blue), random networks (green), and homogeneously mixing populations (red) ($n = 48$, single trial from each condition). In the initial rounds of the game (A–C), the fraction of the population using each name scales as a broad distribution with the overall rank of the name. The scaling is similar in all networks. However, (D) by round 16 in the homogeneously mixing population, a single name breaks the symmetry with the others, accelerating its growth and driving the other options toward extinction (E and F). Over the same time interval, the spatial network and random network exhibit much more moderate growth dynamics, producing exponential distributions with a few competing dominant groups, but without any “winner” emerging. The dashed line in all panels provides a visual guide for the Zipf distribution (slope -1). Note that for $t = 24$ the curve for random networks is absent due to a shorter experimental time.

three network conditions, represented as frequency-rank plots. Early in the evolutionary process, all networks exhibited a broad distribution of active alternatives, indicating that even the least popular options had nontrivial representation within the population. However, as the ecologies evolved, the distribution of alternatives in the spatially embedded and randomly connected populations became increasingly exponential, producing an emergent “oligopoly,” in which a few entrenched local conventions eliminated all other alternatives (45, 52). Each of these conventions competed for ground against the others, but none of them assumed the majority. The ecology evolved quite differently in the homogeneously mixing populations. After the initial

transient, a dominant convention rapidly emerged, breaking the symmetry with its competitors and shifting the population into a “winner take all” regime (27). Despite a large number of competing alternatives circulating in the population (*Robustness*), in every trial in the homogeneously mixing networks, the dynamics converged on a global convention.

More generally, Fig. 4 shows all of the replications of our study. Consistent results were found for each of the topologies at both $n = 24$ and $n = 48$. As a final test of increasing scale, we replicated a trial of the homogeneously mixing population in which the network size was again doubled (27, 33, 53). Fig. 4 shows that a shared social convention spontaneously emerged in a population of

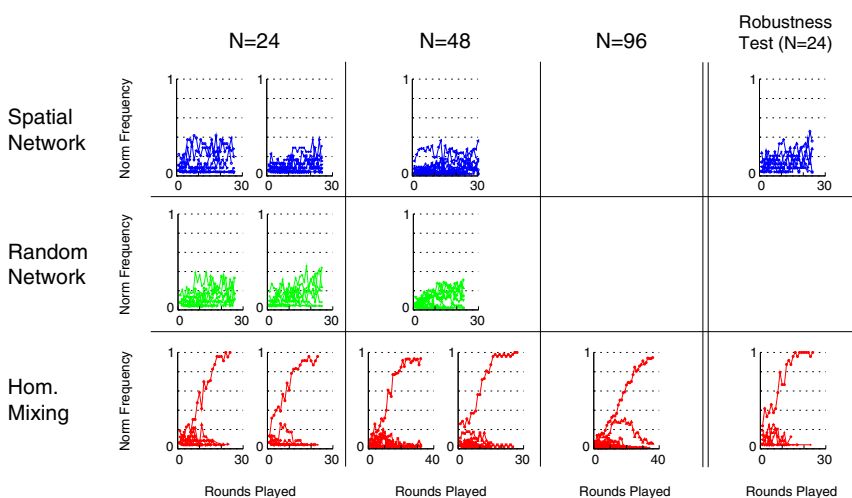


Fig. 4. Size of dominant convention across conditions. The temporal evolution of the ecology of norms is reported for 13 experimental trials. Results show a significant difference in dominant convention size between locally connected networks (spatial and random networks, combined) and homogeneously mixing populations (Wilcoxon Rank Sum Test, $P < 0.01$). In the spatial network and random network trials, the most popular convention was never used by more than 45% of the group and in most trials was well below 40%. Homogeneously mixing populations, on the other hand, produced global social conventions that reached up to 100% of the population. More generally, in spatial networks the average size of the dominant convention across trials was 30% of the population (including $n = 24$ and $n = 48$), in random networks the average size was 33% (including $n = 24$ and $n = 48$), and in homogeneously mixing populations the average size was 96% (including $n = 24$, $n = 48$, and $n = 96$).

$n = 96$ subjects. Global coordination in the $n = 96$ population emerged on a time scale comparable to that of the initial trials ($n = 24$), despite the fact that subjects had no information about how large the coordinating population was. Within the brief time scale of the experimental observations (30 rounds of play on average), large homogeneously mixing populations were significantly more likely ($P < 0.01$) to spontaneously create social conventions than smaller populations with less connectivity.

Discussion

To ensure that our findings do not rely in any way on participants' knowledge of the size of the population or number of interaction partners, we tested the effectiveness of the informational controls used in our experimental design by providing subjects with a postexperiment survey asking them to report (i) the most popular name in their game, (ii) the number of people in their game, and (iii) the number of people with whom they interacted. Across all network conditions and all network sizes, there were no significant differences in subjects' responses regarding the size of their network or the number of neighbors with whom they interacted (see Fig. S5). The only difference in responses was that in all of the homogeneously mixing networks every respondent knew the norm, even though none of them knew how many people were using it.

We also evaluated the robustness of our results for possible biases in the initial distribution of conventions based on external focal points (54). To rule out the possibility that convergence may be biased by the preexisting popularity of some names, we conducted controlled experiments in which participants chose their options from among an arbitrary list of 10 names whose order was randomized at the beginning of the experiment (for each participant, to avoid implicit ranking effects). Results from these controlled experiments are indistinguishable from the results presented above (Fig. 4). Moreover, our findings are further supported by the observed levels of diversity in the emergent ecology of names in each of the other trials. In every trial of our study, the number of suggested names was larger than the size of the population (sometimes by more than a factor of 2; Fig. S4), suggesting that there were no preferred options that initially limited the set of choices in the social evolutionary process.

In sum, our findings demonstrate that social conventions can spontaneously evolve in large human populations without any institutional mechanisms to facilitate the process. Further, the results highlight the causal role played by network connectivity in the dynamics of establishing shared norms. These results contrast with prior work analyzing the effect of network structure on the speed of convergence (6, 41, 55–57). However, that work focuses on the situation where there are just two competing norms that

have different payoff consequences. In our case, by contrast, the number of possible norms is not fixed in advance, and they all have identical payoff consequences. In this case, we find that increased network connectivity can accelerate the rate of convergence to a global norm. As a result, large populations without global information or incentives for collective agreement may nevertheless rapidly self-organize to produce universally shared collective beliefs and behaviors.

We anticipate that our results will be of interest to researchers investigating the effects of online connectedness on the emergence of new political, social, and economic behaviors (58). In particular, a topic of interest for future work will be to explore the practical implications of the unintended effects of increasing social connectedness on the homogenization of behaviors and beliefs among large numbers of individuals who do not even know that they are implicitly coordinating with one another.

Materials and Methods

Participants in the study were recruited at large from the World Wide Web. When participants arrived to play a game, they were randomly assigned to an experimental condition (i.e., a social network) and then randomly assigned to a position within that social network. In a given round of the game, two network neighbors were chosen at random to play with one another. Both players simultaneously assigned names to a pictured object (e.g., a human face), blindly attempting to coordinate in the real-time exchange of naming choices (Fig. S2). If the players coordinated on a name, they were rewarded with a successful payment (\$0.50); if they failed, they were penalized (–\$0.25). (Participants could not go into debt, so failures did not incur a penalty if a participant had a balance of \$0.) After a single round, the participants could see only the choices that they and their partner had made, and their cumulative pay was updated accordingly. They were then randomly assigned to play with a new neighbor in their social network, and a new round would begin. The object that participants were trying to name was the same for the entire duration of the game and for all members of the game. An experimental trial concluded when all members completed the specified number of rounds. Participants did not have any information about the size of the population nor about the number of neighbors to whom they were connected nor even about which individuals they were interacting with in a given round. We explored the dynamics of convention formation over different network sizes ($24 \leq n \leq 96$) and degrees of social connectedness ($4 \leq Z \leq N - 1$). However, the controls within the experiment design ensured that the informational resources provided to subjects were identical across all conditions of the study.

ACKNOWLEDGMENTS. We thank H. P. Young, A. van de Rijt, L. Dall'Asta, A. Kandler, N. Perra, R. Reagans, R. Fernandez, S. Strogatz, and B. Ribeiro for helpful comments and discussion; and A. Wagner, R. Overbey, and R. Chebelev for website development. A.B. thanks A. Vespignani and the MoBS Laboratory at Northeastern University for support in the early stages of the study, and acknowledges support from the Research Pump-Priming Fund of City University London. This work was supported in part by a James S. McDonnell Foundation grant (to D.C.).

- Lass R (1997) *Historical Linguistics and Language Change* (Cambridge Univ Press, Cambridge UK).
- Lewis D (1969) *Convention: A Philosophical Study* (Blackwell, Oxford UK).
- Young HP (1993) The evolution of conventions. *Econometrica* 61:57–84.
- Barkun M (1968) *Law Without Sanctions: Order in Primitive Societies and the World Community* (Yale Univ Press, New Haven, CT).
- Ehrlich PR, Levin SA (2005) The evolution of norms. *PLoS Biol* 3(6):e194.
- Young HP (1998) *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* (Princeton Univ Press, Princeton, NJ).
- Hechter M, Opp K-D (2001) *Social Norms* (Russell Sage Foundation, New York, NY).
- Zhang L, Zhao J, Xu K (2014) Who create trends in online social media: The crowd or opinion leaders? arXiv:1409.0210.
- Boyd R (1988) *Culture and the Evolutionary Process* (University of Chicago Press, Chicago).
- Skyrms B (1996) *Evolution of the Social Contract* (Cambridge Univ Press, Cambridge UK).
- Ullmann-Margalit E (1977) *The Emergence of Norms* (Clarendon Press, Oxford).
- Steels L, Belpaeme T (2005) Coordinating perceptually grounded categories through language: A case study for colour. *Behav Brain Sci* 28(4):469–489, discussion 489–529.
- Correll SJ, Ridgeway CL (2006) *Expectation States Theory. Handbook of Social Psychology*, ed DeLamater J (Springer, Heidelberg), pp 29–51.
- Merton RK (1938) Science and the social order. *Philos Sci* 5(3):321–337.
- Kearns M, Judd S, Tan J, Wortman J (2009) Behavioral experiments on biased voting in networks. *Proc Natl Acad Sci USA* 106(5):1347–1352.
- Young MW (1971) *Fighting with Food. Leadership, Values and Social Control in a Massim Society* (Cambridge Univ Press, Cambridge UK).
- Salganik MJ, Dodds PS, Watts DJ (2006) Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311(5762):854–856.
- Bicchieri C (2006) *The Grammar of Society: The Nature and Dynamics of Social Norms* (Cambridge Univ Press, Cambridge UK).
- Hayek FA (1960) *The Constitution of Liberty* (Routledge).
- Sugden R (1989) Spontaneous order. *J Econ Perspect* 3:85–97.
- Haken H (1988) *Information and Self-Organization: A Macroscopic Approach to Complex Systems* (Springer, Heidelberg).
- Strogatz S (2003) *Sync: The Emerging Science of Spontaneous Order* (Hyperion).
- Steels L (1995) A self-organizing spatial vocabulary. *Artif Life* 2(3):319–332.
- Judd S, Kearns M, Vorobeychik Y (2010) Behavioral dynamics and influence in networked coloring and consensus. *Proc Natl Acad Sci USA* 107(34):14978–14982.
- Young HP (1998) Social norms and economic welfare. *Eur Econ Rev* 42(3):821–830.
- Helbing D, Yu W, Opp K-D, Rauhut H (2014) Conditions for the emergence of shared norms in populations with incompatible preferences. *PLoS ONE* 9(8):e104207.
- Baronchelli A, Felici M, Loreto V, Caglioti E, Steels L (2006) Sharp transition towards shared vocabularies in multi-agent systems. *J Stat Mech* P06014.
- Durrett R, Levin SA (2005) Can stable social groups be maintained by homophilous imitation alone? *J Econ Behav Organ* 57(3):267–286.
- Galantucci B (2005) An experimental study of the emergence of human communication systems. *Cogn Sci* 29(5):737–767.

