Supporting Information

Centola and Baronchelli 10.1073/pnas.1418838112

Experimental Design

Each trial of the study consisted of a network structure, a prespecified number of rounds, and a set of participants equal to the size of the network (N). Participants were randomly assigned to a trial (i.e., an active network), and then they were randomly assigned to a node within that network (Fig. S1). Participants had no information about who or how many individuals they were directly connected to or how many people were in the population. The subject experience was identical in every network condition. Consequently, any differences in collective behavior across network conditions were due to the structure of the interaction networks, and not to information the subjects had about the structure or size of the population in which they were.

Subject Experience

Each round subjects were randomly paired with one of their network neighbors and shown a picture for which they had to enter a name (Fig. S2*A*). They were given a 20-s time limit in which they could enter a name for the pictured face. During the same time interval, their partner was given the same face and the same time limit. If either subject did not complete an answer in the allotted time, the system registered a void answer and that round was considered a "null" round in which no information was exchanged. In terms of the participants' scores, both participants were registered as a failed interaction for a given round if one or both of them produced a void answer. Alternatively, if both players entered a name for the pictured face within the time limit, the round concluded with a page that told both subjects the name their partner entered and indicated whether they successfully matched or not. Their score for that round was indicated as either a "match" or a "no match," respectively, depending on whether they were successful or not. Accordingly, the players' scores were also updated based on whether they succeeded or not (Fig. S2*B*). The players then waited to be assigned with a new partner (another one of their network neighbors). Once a player was assigned, the player was again presented with same game screen and the opportunity to name the image, within the 20-s time interval (Fig. S2*C*). This procedure was repeated for the allotted number of rounds until the player had completed the prespecified number of rounds.

Each player was permitted to move at their own pace through the game. Some players may therefore have completed their allotted number of rounds before some of their network neighbors. To ensure that none of these neighbors was "stranded," with all of their neighbors finishing before they could complete the prespecified number of rounds, some players were given the opportunity to play additional rounds of the game until the slower player had completed his or her full allotment of rounds. These additional rounds were identical to the earlier rounds and ensured that by the end of the game every subject had played at least the prespecified number of rounds. When all subjects had completed the prespecified number of rounds, the game ended.

Subject Recruitment

Participants in our study were recruited at large from the World Wide Web to be players in the "Name Game." When they arrived to the study, each participant completed the registration by choosing a username and an avatar. Participants were then provided with a specific time to return to the site when they would play the live Name Game with a group of anonymous, randomly selected subjects. The study was run for a 140-d period, over which time recruitment campaigns were conducted to attract subjects to participate in the study. In total, 510 subjects were recruited to participate in the study. Of them, 120 participated in spatially embedded lattice networks, 96 participated in randomly connected graphs, and 264 participated in homogeneously mixing networks, as detailed in *Data Analysis* and *Replication*. Subjects were recruited through email advertisements sent to a broad list of websites that subscribe to the Adweek mailing list.

Network Structures

Each trial of the study consisted of a network structure and a population that filled it. We explored three different topologies within our design: (*i*) a spatially embedded network, which was structured as a one-dimensional lattice with degree 4 (each node was tied to nearest neighbors and next-nearest neighbors); (*ii*) a random network topology, which was structured as a random graph with constant degree 4; and (*iii*) a homogeneously mixing population, or a complete graph. Each of these topologies was tested for populations of size $n = 24$ and $n = 48$. For $n = 96$, we conducted a single study with the homogeneously mixing network (*Replication*). As discussed in *Model*, our formal predictions for $n = 96$ were qualitatively the same as those found for $n = 48$ and $n = 24$.

Data Analysis

The experimental data were produced as a chronological sequence of rounds, ordered according to the starting time of the game (Fig. S3). Each round consisted of two participants each typing a name. As described above, users had 20 seconds to type a name, after which time the system registered their answer as void. A round was considered successful only when the two participants entered the same name, irrespective of case. In the analysis, the population's evolutionary time scale is measured in terms of the number of times the entire population goes through a single round of the game, which we refer to as a single "Round Played." A Round Played for the population corresponds to $N/2$ individual rounds in the data sequence. Thus, a Round Played is equivalent to $N/2$ pairs (i.e., the entire population) all playing once. This approach allows us to measure the movement of the entire population through sequential rounds of the Name Game, based on the standard mapping between Monte Carlo steps and microscopic interactions (1). The value of each measured quantity for the population at any Round Played is obtained by averaging over the last $N/2$ individual rounds. Thus, for example, the average success rate of the population at the $(x + 1)$th Round Played corresponds to the average of the individual outcomes in the interval between the lines $(xN/2 + 1)$ and $(x + 1)N/2$ in the data sequence (success being a binary variable is assigned values 0 or 1 at each individual round) (Fig. S3). Analogously, the frequency of each norm at the $(x + 1)$th Round Played is relative to the frequency of the other norms in the same interval. If either member of a given pair fails to type a name within an interaction, then this interaction is treated as void. This is handled differently for the analysis of success versus the analysis of norm ecology. For success rates, a void entry from either member of a pair indicates nonparticipation (i.e., a null interaction) in that round. Thus, in the evaluation of the temporal evolution of the success rate, we disregarded all those interactions in which void entries appear. For analyzing the evolving norm ecology, however, if only one member of a pair produced a void field whereas the other member typed a name, we include the name that was typed by the active participant as a data point in our analysis of the active names within the population. As a consequence, for a single experimental realization, the considered temporal sequence can be slightly different for the series of pairwise success and that of the overall norm ecology. It is worth stressing that all results are robust against variations of the specific procedure adopted to take into account the void field and that different ways of dealing with interactions in which one of the participants did not type any guess produce results that are equivalent, and virtually undistinguishable, under any respect.

Replication

The results presented in the main text were consistent across all replications of our experiments. We replicated the experiment eight times for homogeneous mixing populations ($n = 12$, 24, 48, and 96), six times for the spatial networks ($n = 12$, 24, and 48), and three times for the random networks ($n = 24$ and 48). All trials of size $n = 24$ were run for an average of 25 rounds; trials of $n = 48$ were run for an average of 30 rounds, and the trial of $n = 96$ was run for 40 rounds.

The choice of the above-mentioned trials was dictated by our research questions.

As predicted by the model, small n trials ($n = 12$) produced similar dynamics across all experimental conditions, preventing any identification of the effects of network structure on convention formation. Six trials were conducted with $n = 24$. According to our model, this was the minimal population size at which significant differences in the emergent dynamics of local coarsening versus symmetry breaking could be detected. For robustness, each trial was replicated twice in each condition. This allowed identification of the main dynamical differences across conditions and assessment of the validity of the model predictions, as shown in Fig. 1. The replications with $n = 48$ corroborated these results in larger networks and allowed us to obtain data sufficient to construct the distributional analyses presented in Fig. 3.

Because consensus becomes more difficult with increasing scale, the failures in the spatial lattices and random networks indicated that larger n studies with those topologies would yield similar results. Our focus for additional replications was thus the homogeneously mixing population—that is, the only condition in which a global consensus emerges in our experiments. We replicated the $n = 48$ homogenously mixing experiment a second time, and additionally we tested our results in a final, considerably more demanding trial with 96 participants. As a final test of our findings, we also conducted two "Name List" trials, which served as a further check on the robustness of the results (*Robustness*).

Robustness

A possible concern with the design of our study is that the distribution of words entered by subjects would be skewed in favor of a particularly salient name (where saliency could have been due to a vast range of external events/factors, such as the celebrities in the news and so forth), which would drive convergence by artificially reducing the number of options in the population. To check the robustness of our results in a setting that eliminated these concerns, we replicated our study in homogeneous mixing and spatial networks of size $n = 24$, in an environment in which participants could not type their own name entries. Instead of allowing participants to enter their proposed name in a text box, we provided them a fixed list of 10 names. Participants had to name a feminine face and could select in each round one name from the fixed list of Sophia, Emma, Isabella, Olivia, Ava, Emily, Abigail, Mia, Madison, and Elizabeth—corresponding to the most popular baby names for females for 2012 in the United States according to the US Social Security Office (2). The order of these names was randomized at the beginning of the experiment for each participant, to rule out possible ordering biases. Fig. 4 shows that the results are consistent with those obtained in the trials using open field name entry.

As a secondary check on our results, we examined the data to determine whether the list of actual names that were suggested by subjects in any of the trials was artificially limited to a small number of options. As shown in Fig. S4, in each of the conditions with open name fields, the number of names entered by subjects was greater than the size of the population in the game.

Post-trial User Tests

To ensure that the informational controls in our study were effective, we provided subjects from five selected trials with a poststudy questionnaire, asking them to report (*i*) the number of people in their game and (*ii*) the number of people with whom they directly interacted. Fig. S5*A* shows the mean and SD of responses to the number of people in the game (normalized by the number of rounds that subjects played). There were no significant differences in the average responses across network structures and network sizes ($P > 0.2$, Mann–Whitney U test). Similarly, Fig. S5*B* shows the normalized mean and SD of responses for the number of people with whom players believed they interacted. There were no significant differences in the average responses across network structures and network sizes ($P > 0.2$, Mann–Whitney U test).

Limitations of the Study

As with all experiments, the scientific controls that made this study possible also present limitations. Most notably, practical constraints on the number of people that can be recruited to simultaneously participate in an evolving social convention within an experimental environment prevented us from running larger experiments (i.e., $n > 100$). These constraints also limited the duration of our experiments—that is, the number of rounds of play—as our design relied on subjects' sustained behavioral engagement over the entire study. Although these practical constraints limited the size of our empirical study, the correspondence between our model and the experimental data provides guidance for our expectations about how these evolutionary systems behave at larger sizes and longer time scales. As discussed in *Model*, the results from our simulations suggest convergence time in each of the three network scales as a direct function of the topology. For spatial lattices, convergence is expected to scale as $O(N^2)$ Rounds Played, whereas for both the random graph and the homogeneously mixing population, convergence time is expected to scale as $O(N^{0.5})$ Rounds Played. Based on the agreement between our experimental results and the model, we speculate that the dynamics of norm evolution within each network topology will follow the patterns of coarsening (in the spatial lattice and early stages of the random graph) and symmetry breaking (in the homogeneously mixing population and late stages of the random graph), as described below.

Model

The Naming Game model constructs a population of n agents that engage in pairwise interactions to negotiate local coordination and is able to demonstrate the emergence of a global convention among them (3). An example of such a game is that of a population that has to reach consensus on the name for an object, using only local interactions, as in the Name Game experiment.

In the model, each agent has an internal name inventory in which an a priori unlimited number of words can be stored. As an initial condition, all inventories are empty. At each time step, a pair of agents is chosen randomly, one playing as "speaker" and the other as "hearer," and interact according to the following rules:

- The speaker randomly selects one of her words (or invents a new word if his or her inventory is empty) and conveys it to the hearer;

- if the hearer's inventory contains such a word, the two agents update their inventories so as to keep only the word involved in the interaction (*success*); and
- otherwise, the hearer adds the word to those already stored in her inventory (*failure*).

The nonequilibrium dynamics of the Naming Game are characterized by three temporal regions: (*i*) initially the words are invented; (*ii*) then they spread throughout the system, inducing a reorganization process of the inventories; and (*iii*) this process eventually triggers the final convergence toward the global consensus (all agents possess the same unique word). The dynamics leading to final consensus and the associated scaling of the consensus time with the population size depend crucially on the topological properties of the social network identifying the set of possible interactions among individuals.

Homogeneous Mixing Populations. In homogeneously mixing populations, the third step above is triggered by a symmetry breaking in the ecology of conventions, in which the most popular norm will progressively eliminate all of the competitors. For a population of size N, consensus is reached in a time $t_{conv} \sim 0(N^{1.5})$ microscopic interactions—that is, in $O(N^{0.5})$ Rounds Played—according to our definition (3).

The symmetry-breaking process has been clarified analytically by considering a system prepared in an initial configuration in which half of the population knows only convention "A" and the other half only convention "B" (3, 4). Here, stochastic fluctuations break the initial symmetry between the two norms, making one of them more popular, and the interaction dynamics amplify this small advantage until a final state in which the initially disadvantaged convention is extinct. Thus, in the limit of large population, any initial imbalance in favor of one of the two conventions will eventually determine the success of that convention (4, 5).

The situation is more complex when more than two conventions are present in the system, but the overall symmetry-breaking picture remains the same (3, 5). This mechanism is radically different from what is observed in pure imitation models, such as the Moran process or the voter model, where fluctuations dominate the whole of the process, leading to consensus, and the advantage of one of the competing states can be reversed easily during the dynamics of the process (6). When only pure imitation is at work, the fluctuation- driven consensus is reached in $O(N)$ Rounds Played (or Monte Carlo steps).

Networked Populations. In the model described above, at each time step two agents are randomly selected. The assumption behind this homogeneous mixing, or "mean-field," rule is that the population is not structured and that any agent can in principle interact with any other. However, when actors are embedded in a fixed network, the topology in which the population is embedded identifies the set of possible interactions among the individuals. Thus, the group of communicating individuals can be described as a network in which each node represents an agent and the links connecting different nodes determine the allowed communication channels. The (statistical) properties of the underlying network significantly affect the overall dynamics of the model.

Lattices. On low-dimensional lattices, each agent can rapidly interact two or more times with its neighbors, favoring the establishment of a local consensus with a high success rate—that is, of small sets of neighboring agents sharing a common unique word. As the process evolves, these "clusters" of neighboring agents with a common unique word undergo a coarsening phenomenon with competition among them driven by the fluctuations of the interfaces. The coarsening picture can be extended to higher dimensions, and the scaling of the convergence time has been shown to be $O(N^{2/d})$, where $d \leq 4$ is the dimensionality of the space (7). This prediction has been confirmed numerically.

Small-world networks. Results concerning the homogeneously mixed population, on the one hand, and regular lattices, on the other, act as fundamental references to understand the role of the different properties of complex networks. In between these regimes, the small-world network (8) allows us to interpolate progressively from regular structures to random graphs by tuning the p parameter describing the probability that a link of the regular structure is rewired to a random destination. The main result is that the presence of shortcuts, linking agents otherwise far from each other, allows recovering the fast convergence typical of the mean-field case. The finite connectivity, on the other hand, guarantees that there will be a good degree of coordination between neighbors from the start of the dynamics, as in regular structures.

In these randomized topologies, two different regimes are observed (9). For times shorter than a cross-over time, $t_{cross} = O(N/p^2)$, one observes the usual coarsening phenomena, as the clusters are typically one-dimensional—that is, because the typical cluster size is smaller than $1/p$. For times much larger than t_{cross}, the dynamics shift. They become dominated by the existence of shortcuts and follow the mean-field behavior similar to the one observed on the complete graph. The convergence time, measured in microscopic interactions, scales therefore as $N^{1/2}$ and not as $N^{2/d}$ (as in low-dimensional lattices) (9).

Complex networks. Most of the relevant features exhibited by complex networks have been explored systematically, mainly by means of computer simulations. The scaling exponents observed in both homogeneous [e.g., Erdös–Rényi (10)] and heterogeneous [e.g., Barabási–Albert (11)] networks are similar to the one observed in the Watts–Strogatz small-world graphs (8) for both consensus time and memory use. In particular, the scaling laws observed for the convergence time is a general robust feature that is not affected by further topological details, such as the average degree, the clustering, or the particular form of the degree distribution (12).

Robustness of the Model. The model described above has been modified in several directions to test its robustness (6, 13–21). For example, the rule describing how a word is selected from the inventory has been investigated, and more efficient strategies have been identified (13). In the same way, the symmetric update of the inventories has been altered, and the role of the feedback between the agents has been investigated (16, 21). However, all permutations of the model exhibit qualitatively similar dynamics, which rely on two essential elements: (*i*) memory and (*ii*) the fact that for a success to take place both agents must have already heard the successful convention in the past. That is, actors depend upon multiple exposures to a term to successfully coordinate on it (22). These two elements seem to be crucial to reproduce the observed dynamics of coarsening on low-dimensional lattices and symmetry breaking in random networks and homogeneously mixing populations.

Model Rules and User Behavior

To test the relationship between the model and the experimental results in more depth, we investigated how well the individual behavior of participants in the study matched with the theoretical model. We then simulated the long-term dynamics of the real user behaviors and compared it to the expected dynamics from the theoretical model.

In the theoretical model, agents accrue a list of word options based on their history of interactions. The only words that they can enter at a given round are those that currently exist in their inventory. If they experience a successful match, their inventory is deleted except for the matching word. The inventory can increase again through subsequent interactions, however any subsequent matches again reset the inventory to 1, leaving only the most recent matching word. We evaluated subjects' behavior in terms of whether the answers they provided at every round were consistent with answers that agents in our theoretical model could have provided; that is, we evaluated whether the answers that subjects actually used would have been in their "inventories" if they had followed the same rules as the model, given their histories of past interactions, failures, and successes.

Remarkably, we found a 95% agreement between the model and the subjects' behaviors. In other words, 95% of the time, subjects' choices were entirely consistent with the rules of the theoretical model. When these individual dynamics were simulated (95% model rules, 5% random entries—either through novel word choice or through choosing words from a deleted inventory), the collective dynamics were indistinguishable from those of the theoretical model. Consistent with the model, these results suggest that subjects' behaviors were governed more by their recent successes than by their history of past plays.

Model Implications

As shown in the main text, the model captures well the results of the Name Game experiment, and the microscopic rules provide a good fit with the empirically measured user behavior. However, as discussed above, experimental constraints limit the region of accessible parameters, in particular with respect to the duration of an experiment and the population size. The model behavior allows us to make grounded predictions on the outcome of experiments at larger scales. In particular, we would expect that:

a) The dynamics observed on the random graph would eventually be different from the one of spatial networks, and the scaling of the convergence time with the population size will be similar to the one observed in the homogeneously mixing population (9).

b) The difference between the initial stages of the spatial graph and the homogeneously mixing population will be more and more significant (12).

c) The symmetry breaking transition, which governs the consensus process in the homogeneously mixing population, will result in a characteristic S-shaped behavior of the success rate curve (3).

Fig. S6 shows the evolution of the space of norms for the three networks considered in the main text—(Fig. S6 *A* and *D*) spatial lattice networks, (Fig. S6 *B* and *E*) random graphs, and (Fig. S6 *C* and *F*) homogeneously mixing—for populations of $n = 48$. Fig. S6 *A–C* show the results for the experimentally accessible regions of the dynamics (i.e., 30 rounds). Fig. S6 *D–F* show the same simulations with time scales extended until final convergence. To demonstrate these effects at larger population scales, Fig. S7 shows the evolution of the player success rate until model convergence in all three network conditions for populations of size $n = 1{,}000$.

1. Marro J, Dickman R (1999) *Nonequilibrium Phase Transitions in Lattice Models* (Cambridge Univ Press, Cambridge, UK).
2. 2013) *Top 10 Baby Names for 2012*. Available from the Official Social Security Website, www.ssa.gov/OACT/babynames/.
3. Baronchelli A, Felici M, Loreto V, Caglioti E, Steels L (2006) Sharp transition towards shared vocabularies in multi-agent systems. *J Stat Mech* P06014.
4. Baronchelli A, Loreto V, Steels L (2008) In-depth analysis of the Naming Game dynamics: The homogeneous mixing case. *Int J Mod Phys C* 19(05):785–812.
5. Baronchelli A, Dall'Asta L, Barrat A, Loreto V (2007) Nonequilibrium phase transition in negotiation dynamics. *Phys Rev E Stat Nonlin Soft Matter Phys* 76(5 Pt 1):051102.
6. Castellano C, Fortunato S, Loreto V (2009) Statistical physics of social dynamics. *Rev Mod Phys* 81(2):591–646.
7. Baronchelli A, Dall'Asta L, Barrat A, Loreto V (2006) Topology-induced coarsening in language games. *Phys Rev E Stat Nonlin Soft Matter Phys* 73(1 Pt 2):015102.
8. Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. *Nature* 393(6684):440–442.
9. Dall'Asta L, Baronchelli A, Barrat A, Loreto V (2006) Agreement dynamics on small-world networks. *Europhys Lett* 73(6):969–975.
10. Erdos P, Rényi A (1959) On random graphs. *Publicationes Mathematicae Debrecen* 6: 290–297.
11. Barabási A-L, Albert R (1999) Emergence of scaling in random networks. *Science* 286(5439):509–512.
12. Dall'Asta L, Baronchelli A, Barrat A, Loreto V (2006) Nonequilibrium dynamics of language games on complex networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 74(3 Pt 2):036105.
13. Baronchelli A, Dall'Asta L, Barrat A, Loreto V (2006) Strategies for fast convergence in semiotic dynamics. *Artificial Life X—Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, eds Rocha LM, Yaeger LS, Bedau MA, Floreano D (The MIT Press, Cambridge MA), pp 480–485.
14. Wang W, Lin B, Tang C, Chen G (2007) Agreement dynamics of finite-memory language games on networks. *Eur Phys J B* 60(4):529–536.
15. Dall'Asta L, Castellano C (2007) Effective surface-tension in the noise-reduced voter model. *EPL* 77(6):60005.
16. Lu Q, Korniss G, Szymanski BK (2008) Naming games in two-dimensional and small-world-connected random geometric networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 77(1 Pt 2):016111.
17. Brigatti E (2008) Consequence of reputation in an open-ended naming game. *Phys Rev E Stat Nonlin Soft Matter Phys* 78(4 Pt 2):046108.
18. Brigatti E, Roditi I (2009) Conventions spreading in open-ended systems. *New J Phys* 11(2):023018.
19. Lipowski A, Lipowska D (2009) Language structure in the n -object naming game. *Phys Rev E Stat Nonlin Soft Matter Phys* 80(5 Pt 2):056107.
20. Lei C, Jia J, Wu T, Wang L (2010) Coevolution with weights of names in structured language games. *Physica A* 389(24):5628–5634.
21. Baronchelli A (2011) Role of feedback and broadcasting in the naming game. *Phys Rev E Stat Nonlin Soft Matter Phys* 83(4 Pt 2):046103.
22. Centola D, Macy M (2007) Complex contagions and the weakness of long ties1. *Am J Sociol* 113(3):702–734.

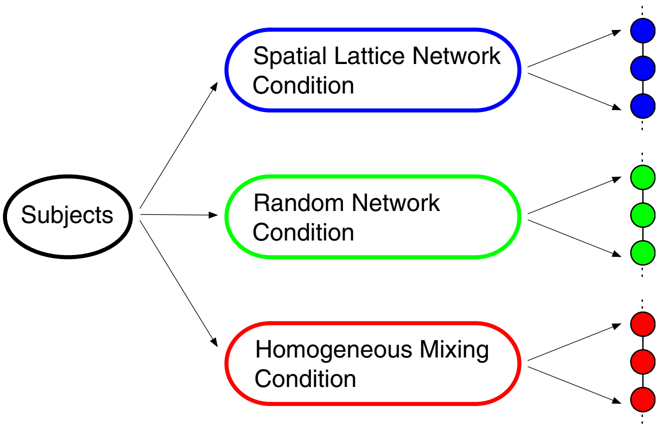

Fig. S1. Schematic representation of randomization to conditions. Subjects arriving to the study were first randomly assigned to an experimental condition (i.e., a social network) and then randomly assigned to a specific node within that network. The nodes directly connected to an individual constituted the set of potential partners that he or she could interact with during the game.

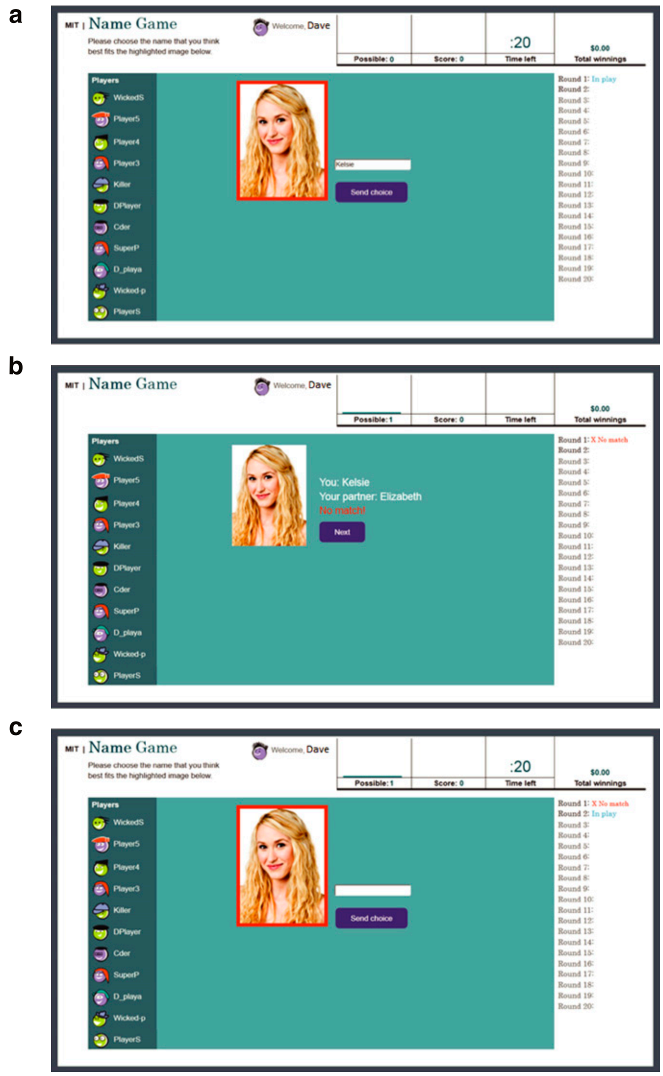

Fig. S2. User interface and experience. At the start of the study, (*A*) subjects are given a picture of a face and an entry field (with no character limit) to provide a name. Once a subject submits his or her choice (*B*) he or she is then exposed to the choice of her interaction partner. If the choices match, both subjects are rewarded; otherwise, there is no reward for that round. The study then advances to the next round (*C*), where the player is again matched with a partner and repeats the same procedure. Each round lasts for 20 s maximum, and the players have real-time information on their record of matches and failures over the past rounds. The "Players" column on the left of the screen is a static representation of other player icons, identical across experimental conditions (and hence independent from the actual topology and population size).

Round	Player 1	Name 1	Player 2	Name 2	Success	Round Played	Player Success Rate
1	A	"Sarah"	L	"Isabella"	0		
2	F	"Maria"	B	"Anna"	0		
3	G	"Isabella"	I	"Isabella"	1		
4	D	"Lauren"	C	"Sarah"	0		
5	E	"Giulia"	A	"Anne"	0	1	1/5
6	F	"Mary"	H	"Mary"	1		
...	...	...	...	...	...		

Fig. S3. Schematic representation of the data for a population of size $n = 10$. Every experiment generates an ordered list of individual rounds, ordered on the basis of their starting time. If the names typed by the two players are the same, the interaction is a success, and the success variable takes a value of 1; otherwise, it is a failure and the relative variable is set at 0. Global quantities, such as the Player Success Rate in the figure, are averaged over $N/2$ individual rounds, corresponding to one Round Played. In the figure, the Player Success Rate in the first Round Played success is equal to 1/5, as one pair out of five achieved success. Within each Round Played, each player plays once on average. For instance, in Round Played 1, user A plays twice (rounds 1 and 5) and player H does not play at all, whereas in Round Played 2, it would be reversed.

	Cumulative Number of Alternatives Created In Trial
Trial 1 (N=24, Spatial Network)	53
Trial 2 (N=24, Spatial Network)	83
Trial 3 (N=24, Random Network)	40
Trial 4 (N=24, Random Network)	30
Trial 5 (N=24, Homophilous Mixing)	50
Trial 6 (N=24, Homophilous Mixing)	48
Mean	50.66
Standard Deviation	17.88
Trial 7 (N=48, Spatial Network)	66
Trial 8 (N=48, Random Network)	51
Trial 9 (N=48, Homophilous Mixing)	64
Trial 10 (N=48, Homophilous Mixing)	52
Mean	58.25
Standard Deviation	7.84
Trial 11 (N=96, Homophilous Mixing)	120
Mean	-
Standard Deviation	-

Fig. S4. Cumulative names entered over the course of the study. The numbers reported here indicate the number of different words in active circulation in each population. Identical spellings with different cases were considered to be the same word. The number of alternative names created in each trial was larger than the population size.

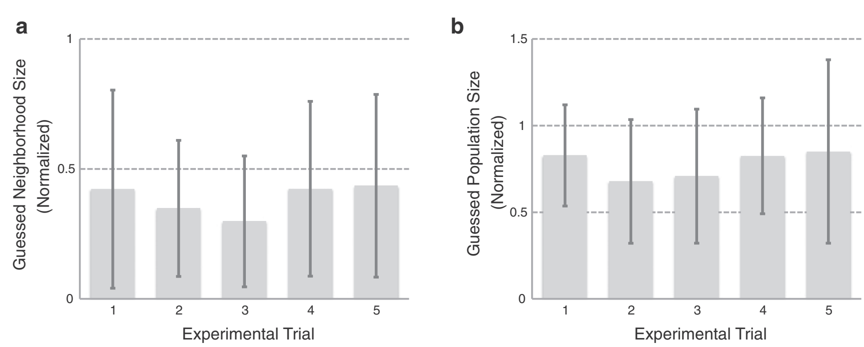

Fig. S5. Subjects' informational awareness about experimental conditions. (*A*) Subjects' reported beliefs about the number of people they interacted with are shown for five experimental trials (normalized by the number of rounds that subjects played). Results are shown for a representative (*1*) Spatial Network $n = 24$, (*2*) Homogeneous Mixing $n = 24$, (*3*) Random Network $n = 48$, (*4*) Homogeneous Mixing $n = 48$, and (*5*) Homogeneous Mixing $n = 96$. For the same trials, *B* shows subjects' reported beliefs about the population size of their study (normalized by the number of rounds that subjects played). Error bars indicate the SD in responses.

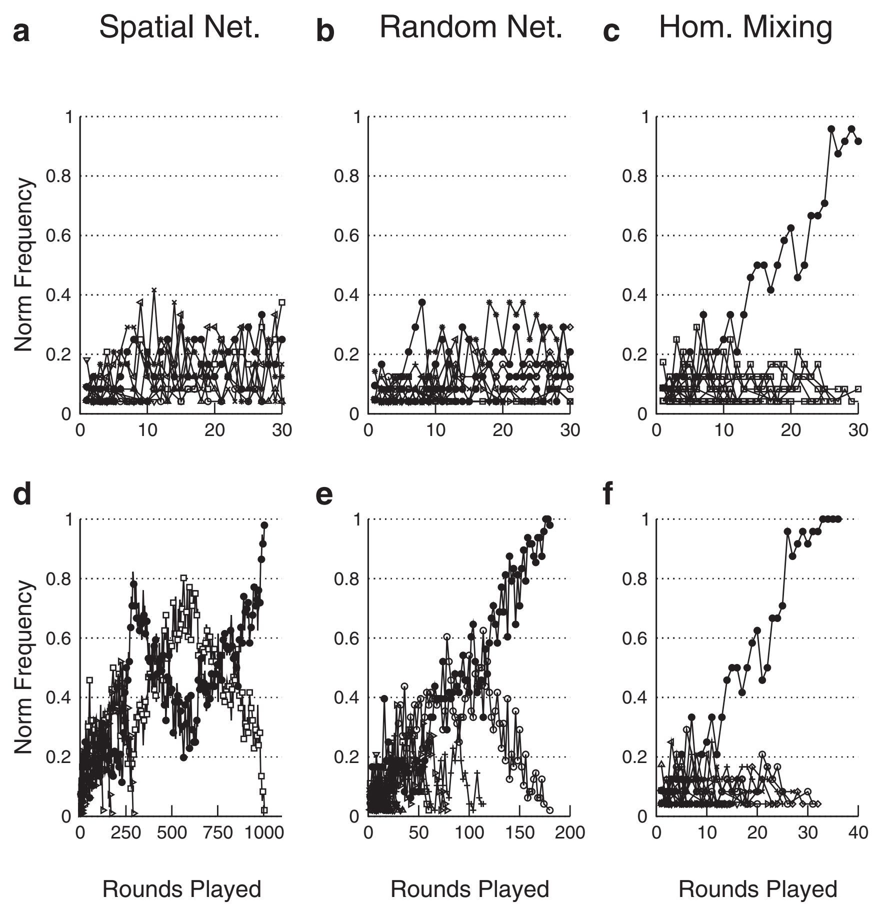

Fig. S6. Numerical simulations of evolution to final consensus. (*A* and *D*) Spatial network, (*B* and *E*) random network, (*C* and *F*) homogeneously mixing population. The top panels (*A*–*C*) show the temporal evolution in the first 30 rounds—that is, in the experimentally accessible regions of the dynamics. The bottom panels (*D*–*F*) show the same simulations run until final convergence. For the spatial network, final consensus requires more than 1000 rounds of play. After ~300 rounds, only two conventions remain in the population, and they swap their rank twice as the dynamics proceeds through local coarsening. In the random graph, the initial clusters of local coordination are more permeable due to the lack of a spatial structure and therefore permit symmetry breaking on longer time scales. In the figure, local coarsening shifts to symmetry breaking within ~120 rounds, and the population reaches convergence by 180 rounds. In the homogeneously mixing population, rapid symmetry breaking leads to convergence on accelerated time scales detectable within the experimental window.

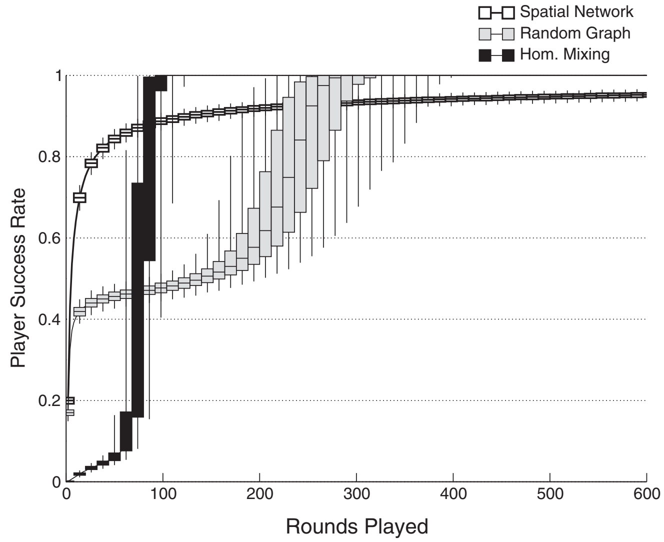

Fig. S7. Large *n* simulations of the Player Success Rate ($n = 1,000$). Average success in the spatial networks (white boxes) grows rapidly in the first 50 rounds because repeated interactions within neighborhoods facilitate local coordination. However, final consensus time is protracted by local competition among emergent groups. In random networks (gray boxes), an initial phase of rapid local consensus with slow global consensus shifts toward a sudden jump of the Player Success Rate. This is a typical signature of the underlying symmetry breaking process eventually occurring in the space of conventions. In homogeneously mixing populations (black boxes), initial coordination is most difficult during the first 50 rounds. However, this gives way to a sharp symmetry breaking transition, in which a global convention spontaneously emerges.