

Heterogeneous networks do not promote cooperation when humans play a Prisoner's Dilemma

Carlos Gracia-Lázaro^a, Alfredo Ferrer^a, Gonzalo Ruiz^a, Alfonso Tarancón^{a,b}, José A. Cuesta^{a,c}, Angel Sánchez^{a,c,1}, and Yamir Moreno^{a,b,1}

^aInstituto de Biocomputación y Física de Sistemas Complejos (BIFI), Universidad de Zaragoza, 50018 Zaragoza, Spain; ^bDepartamento de Física Teórica, Universidad de Zaragoza, 50009 Zaragoza, Spain; and ^cGrupo Interdisciplinar de Sistemas Complejos (GISC), Departamento de Matemáticas, Universidad Carlos III de Madrid, 28911 Leganés, Madrid, Spain

Edited by Simon A. Levin, Princeton University, Princeton, NJ, and approved June 8, 2012 (received for review April 24, 2012)

It is not fully understood why we cooperate with strangers on a daily basis. In an increasingly global world, where interaction networks and relationships between individuals are becoming more complex, different hypotheses have been put forward to explain the foundations of human cooperation on a large scale and to account for the true motivations that are behind this phenomenon. In this context, population structure has been suggested to foster cooperation in social dilemmas, but theoretical studies of this mechanism have yielded contradictory results so far; additionally, the issue lacks a proper experimental test in large systems. We have performed the largest experiments to date with humans playing a spatial Prisoner's Dilemma on a lattice and a scale-free network (1,229 subjects). We observed that the level of cooperation reached in both networks is the same, comparable with the level of cooperation of smaller networks or unstructured populations. We have also found that subjects respond to the cooperation that they observe in a reciprocal manner, being more likely to cooperate if, in the previous round, many of their neighbors and themselves did so, which implies that humans do not consider neighbors' payoffs when making their decisions in this dilemma but only their actions. Our results, which are in agreement with recent theoretical predictions based on this behavioral rule, suggest that population structure has little relevance as a cooperation promoter or inhibitor among humans.

evolutionary game dynamics | network reciprocity | conditional cooperation

The strong cooperative attitude of humans defies the paradigm of *Homo economicus* and poses an evolutionary conundrum (1, 2). This conundrum is because many of our interactions can be framed as Prisoner's Dilemmas (3–5) or Public Goods Games (6), famous for bringing about a tragedy of the commons (7). Several mechanisms have been suggested as putative explanations of cooperative behavior (8), among which the existence of an underlying network of contacts constraining who one can interact with has received very much attention. This mechanism was first proposed in the work by Nowak and May (9), where simulations on a square lattice with agents that imitate the behavior of their neighbor with the highest payoff showed high levels of cooperation in the Prisoner's Dilemma. The ensuing two decades have witnessed a wealth of theoretical studies that have concluded that this so-called network reciprocity (8) is, indeed, possible under a variety of circumstances, but in many other contexts, networks do not promote—or they even inhibit—cooperation (10, 11). The effect of regular and homogeneous networks on cooperation is very sensitive to the details of the model (e.g., dynamics and clustering), whereas theoretical results and simulations indicate that heterogeneous networks should be particularly efficient in fostering cooperation in social dilemmas (11–13). A natural way to shed some light on these partially contradictory results would be to test experimentally the predictions of the different models. Such tests are currently lacking (14), because the few available experimental works only deal—

with some exception (15)—with very small networks (16–18). Interestingly, the only theoretical result (19) that takes into account the behavioral information extracted from experiments predicts that neither homogeneous nor heterogeneous networks would influence the cooperative behavior in the Prisoner's Dilemma (i.e., the observed cooperation level should be the same as if every player interacted with every other player).

Here, we close the cycle by testing the above theoretical predictions (19) and contributing to the current debate on the existence and effects of network reciprocity by performing experiments on large samples of structured populations of individuals who interact through a Prisoner's Dilemma (PD) game. Specifically, we have designed a setup in which 1,229 human subjects were placed in either a square lattice or a scale-free network, and for more than 50 rounds, they played a 2 × 2 multiplayer PD game with each of their k neighbors, taking only one action [either to cooperate (C) or defect (D)—the action being the same against all opponents]. The experiment was simultaneously carried out on two different virtual networks: a 25 × 25 lattice with $k = 4$ and periodic boundary conditions (625 subjects) and a heterogeneous network with a fat-tailed degree distribution (604 subjects; the number of neighbors varied between $k = 2$ and $k = 16$). Fig. 1 depicts a snapshot of a visual representation of the experiment as well as the two networks; more details on the experimental setup as well as a summary of the actions of the subjects during the experiment can be found in *SI Materials and Methods* and *SI Results and Discussion*. Subjects played a repeated (weak) PD with all their neighbors for an initially undetermined number of rounds. Payoffs of the PD were set to be 7 Experimental Currency Units (ECUs) for mutual cooperation, 10 ECUs for a defector facing a cooperator, and 0 ECUs for any player facing a defector (weak PD) (9). We note that this choice of payoffs is like the experiment in the work by Grujić et al. (15) on a smaller regular lattice (Fig. 1), and therefore, cooperation should reach a high level according to the available simulations (9, 11–13). The size of each network was large enough, and therefore, clusters of cooperators could form (the underlying mechanism by which cooperators may thrive) (20, 21).

On this general setup, we carried out two treatments, which we will refer to as experiment and control. In the experiment, subjects remained at the same positions in the network with the same neighbors throughout all of the rounds played. In the control

Author contributions: A.T., J.A.C., A.S., and Y.M. designed research; C.G.-L., A.F., G.R., A.T., J.A.C., A.S., and Y.M. performed research; C.G.-L., J.A.C., A.S., and Y.M. analyzed data; A.F. and G.R. designed and were in charge of the experimental platform; and J.A.C., A.S., and Y.M. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

See Commentary on page 12846.

¹To whom correspondence may be addressed. E-mail: anx@math.uc3m.es or yamir.moreno@gmail.com.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1206681109/-DCSupplemental.

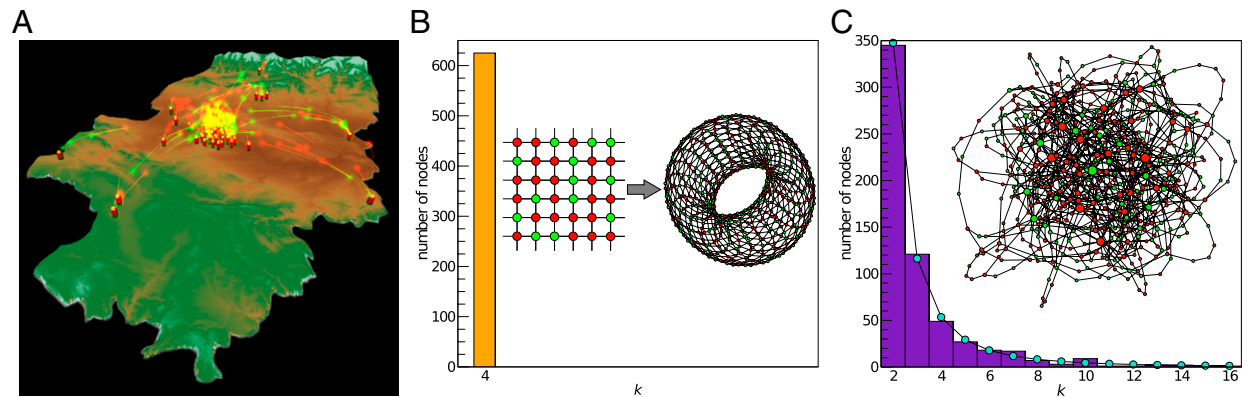


Fig. 1. Players in the experiment were sitting in different physical locations but played in two virtual networks. **A** is a snapshot at round 10 of a graphic animation illustrating the activity during the experiment (*SI Results and Discussion*). On a map of Aragón, the image displays small buildings representing the schools. Arrows (green for cooperate and red for defect) represent actual actions taken by players. They travel to the school, where their randomly assigned neighbors were sitting. Buildings are colored green and red, proportional to the respective number of cooperative and defective actions taken by the subjects in that school. The height of the yellow column on top of each building is proportional to the school's accumulated payoffs. **B** and **C** show snapshots of the two networks at that same round along with their degree distributions (in the case of the heterogeneous network, both the theoretical distribution and the actual realization corresponding to the network of the experiment are represented). Colors indicate the corresponding player's action (green, cooperate; red, defect). The size of a node is proportional to its degree.

treatment, we removed the effect of the network by shuffling the neighbors of each subject in every round. Therefore, in this phase, the players were always connected to the same number of neighbors, but these neighbors changed from round to round. On the screen, subjects saw the actions and normalized payoffs of their neighbors from the previous round, who in the control treatment, were different from their current neighbors with high probability (*SI Materials and Methods* and *SI Results and Discussion*). All treatments of the experiment were carried out in sequence with the same subjects. Players were also fully informed of the different setups that they were going to run. The number of rounds in each treatment was randomly chosen between 50 and 70 to

avoid subjects knowing in advance when it was going to finish, resulting in 51 and 59 rounds for the experimental and control treatments, respectively. Full details are provided in *SI Materials and Methods* and *SI Results and Discussion*.

Results and Discussion

Fig. 2 **A** and **B** shows the fraction of cooperative actions, c , in each round for the two networks and both treatments. The first feature worth noticing in Fig. 2 is that, in the experiment phase, the level of cooperation in either network quickly drops from initial values around 60% to values around 40% and finally settles at a slower pace around 30%, much lower than theoretical models

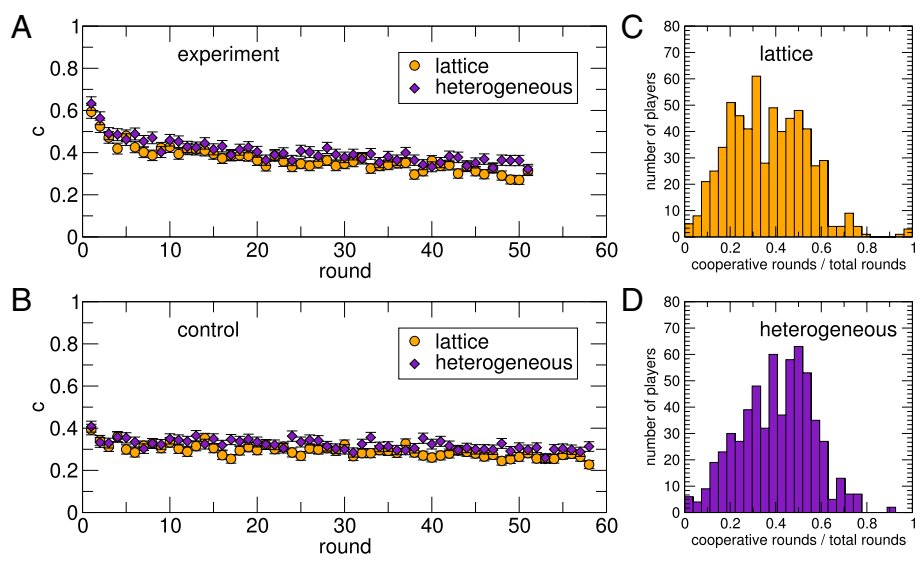


Fig. 2. The level of cooperation declines and is independent of the network of contacts. Fraction of cooperative actions (level of cooperation) per round during the experiment (**A**) and the control (**B**) for both networks and histograms of cooperative actions in the lattice (**C**) and the heterogeneous network (**D**). The histograms (**C** and **D**) show the number of subjects ranked according to the fraction of cooperative actions that they perform along the experiment in the two networks. A Kolmogorov–Smirnov test shows that the distributions are statistically indistinguishable (*SI Results and Discussion*). They illustrate the high heterogeneity in subjects' behavior—their levels of cooperation ranging from nearly zero to almost one in a practically continuous distribution. The corresponding histograms for the control (*Fig. S4*) show that a sizable group of subjects lowered their levels of cooperation, hence becoming mostly defectors. Actually, the decline in the level of cooperation observed in the experiment (**A** and **B**) can be explained as a constant flow of subjects to more defective strategies (evidence supporting this hypothesis in *Figs. S5* and *S6*).

predict (9–11). This finding is especially remarkable for the heterogeneous network, on which no previous results are available, and it is in stark contrast with the predictions that this kind of networks should be particularly efficient in promoting cooperation (11–13). In the control, the initial level of cooperation is already at these low values. This behavior is consistent with previous findings in experiments with smaller lattices (15, 18) as well as unstructured populations (22, 23). Regarding the slow decay undergone by these curves after the first quick drop in the level of cooperation, we believe that this finding is associated with a process of learning (see below). However, the most remarkable result that Fig. 2 provides is that, quite unexpectedly, the network does not have any influence in the evolution of the level of cooperation. In fact, both curves are nearly identical—the slightly lower values obtained for the lattice are likely to arise from the small difference in the initial level of cooperation—despite the very different nature of the networks of contacts between the players.

The experimental result that we have just reported is in very good agreement with the theoretical prediction in ref. 19. This finding prompts us to investigate in detail the players' behavior, because the reason why this prediction was different from earlier ones is the use of the update rule observed in ref. 15. The distributions of subjects by their individual cooperation levels (averaged over the whole experiment) depicted in Fig. 2 *C* and *D* show some heterogeneity of behavior: a few subjects have a high level of cooperation (above 70%), and a sizable fraction cooperated in less than 20% of the rounds, whereas the bulk of subjects have intermediate levels of cooperation. Importantly, the comparison of these distributions of actions, which turns out to be statistically indistinguishable (Kolmogorov–Smirnov test data in Table S1), provides additional evidence that the behavior observed in the two networks is the same. This finding, along with the identical behavior of the cooperation level, suggests that subjects use the same strategies in the lattice and the heterogeneous network, regardless of the fact that, in the latter, the number of neighbors of each individual is heterogeneously distributed.

After considering the aggregate distribution of actions, let us now look for deeper insights on the individual behaviors. As in previous experiments on smaller lattices (15, 18) or unstructured populations (22, 23), our results are compatible with a coexistence of at least three basic strategies: cooperators (players who cooperate with a high probability regardless of the context), defectors (players who defect with a high probability regardless of the context), and moody conditional cooperators (15) (players whose action depends on their previous action as well as the level of cooperation in their neighborhood). A search for moody conditional cooperation shows the results depicted in Fig. 3. Fig. 3 *A* and *B* shows the fraction of cooperative actions occurred after a cooperation/defection as a function of the level of cooperation in the neighborhood. The plots are the fingerprint of moody conditional cooperation: players are more prone to cooperate the more their neighbors cooperate if they cooperated than if they defected. Furthermore, Fig. 3 also supports the striking finding that the strategic behavior of subjects is remarkably similar whether they are playing on the lattice (Fig. 3*A*) or the heterogeneous network (Fig. 3*B*). However, Fig. 3 *C* and *D* shows that the next action of a subject cannot be predicted knowing the largest payoff difference that he/she sees in the neighborhood, thus confirming that subjects did not use payoff differences as a guidance to update their actions.

Fig. 4 provides additional evidence of the significance of the moody conditional cooperation by means of a nonparametric bootstrap check. The series of actions taken by every individual is randomly reassigned to other positions in the lattice or the network, and the probability of cooperation is recomputed. This action is done 10^6 times, and the results show that the two probabilities become independent of the context. Of course, such a reshuffling will not change the dependence on the player's own

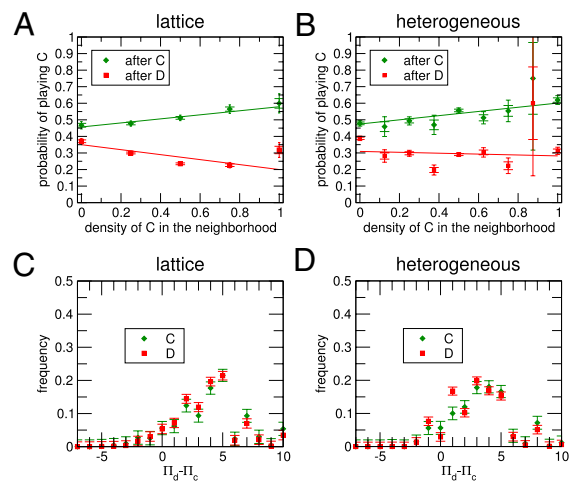


Fig. 3. Players' behavior depends both on the level of cooperation in the neighborhood and their previous action. Frequency of cooperative actions after a cooperative/defective action conditioned to the context (fraction of cooperative actions in the neighborhood in the previous round) observed in the lattice (*A*) and the heterogeneous network (*B*). Details of the linear fits and comparison with randomizations to prove statistical significance can be found in Table S2 and *SI Results and Discussion*. The plots show that there is a relevant dependence on the context for subjects who cooperated in the previous round (i.e., were in a cooperative mood): the cooperation probability increased with the fraction of cooperative neighbors, similar to the conditional cooperators found in the work by Fischbacher et al. (24). However, after having defected, this dependence is less clear, and if anything, it suggests an exploiting behavior—subjects who defected are less prone to cooperate the more cooperation they find. *C* and *D* show how subjects who cooperated or defected are distributed according to the largest payoff-per-link difference in their neighborhoods between the two actions. These plots reveal that a player's decision to cooperate or defect was independent of the payoffs per link that they observed (information that was explicitly provided during the experiment).

previous action, because the order of the actions is not altered; consequently, there are still two distinct lines corresponding to the probability of cooperation after cooperation or defection, but the dependence on the number of cooperators in the previous round is fully removed.

The existence of (almost pure) cooperators and defectors aside from moody conditional cooperators can be further supported through a comparison with the same histograms but for the control condition (Fig. S4), because for the latter, a larger number of subjects are in the region that would correspond to defectors. This finding can be interpreted as an indication that a fraction of (probably) moody conditional cooperators changed to a defective strategy, given that retaliation is ineffective in the control condition. Furthermore, performing running averages of the levels of cooperation during the experiment condition (Figs. S5 and S6) shows that the number of subjects with levels of cooperation that are below a given threshold increases with time—irrespective of the precise value of the threshold. Not only does this finding give support to the existence of this kind of players, but it is consistent with a continuous (albeit small) flow of players who change from moody conditional cooperation to defection—a behavior that could be understood as a generalized form of a grim strategy. Notice that this flow can account for the slow decay observed all along the run of the experiment and control observed in Fig. 2 *A* and *B*.

Finally, another important point that our experiment addresses to some extent is the dependence of the actions on the connectivity of the participants for the heterogeneous network. The results are displayed in Fig. 5, where we represent the average cooperation level c as a function of the connectivity of the players

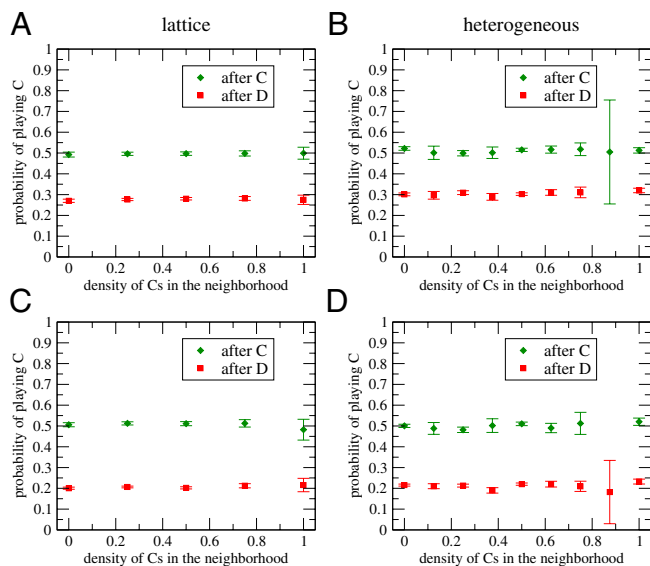


Fig. 4. Null hypothesis statistical significance test. Probability of cooperating after playing C or D, conditioned to the context (fraction of cooperative actions in the neighborhood in the previous round), averaged over 10^6 random shuffling of players. *A* corresponds to the experimental treatment in the lattice, *B* corresponds to the same treatment for the heterogeneous network, *C* corresponds to the control phase in the lattice, and *D* corresponds to the same control treatment for the heterogeneous network. The results show that there is no dependence on the context and hence, that the results of Fig. 3 *A* and *B* are statistically relevant. The anomalous variance (or even absence of data) observed at a fraction of Cs in the neighborhood close to 0.9 is not a relevant feature of the experimental results but a consequence of the very low probability of having events contributing to that bin of the histogram in the heterogeneous network. This anomaly can also be noticed in Fig. 3.

k for both treatments: experiment and control. As can be seen from the plots, there might be some trend to lower levels of cooperation with increasing degree for small connectivities, particularly in the control (the levels for the first three values of the degree in the experiment are not statistically different). However, looking at Fig. 5 as a whole, it becomes clear that there does not seem to be any statistically significant trend. It has to be borne in mind that, in this type of network, the number of hubs or large-degree nodes is intrinsically small, and therefore, the statistics for them are not very accurate (notice the size of the error bars). Going beyond these results would require much larger networks (which would still have the same problem for their higher-degree nodes). Additionally, Fig. 5, *Lower* shows the frequency of cooperative actions of nodes with degree k after playing as C or D with respect to the fraction of their neighbors that cooperated in the previous round. The results are clear evidence that moody conditional cooperation is, indeed, the general behavior, even if one disaggregates the data in terms of their degree. As we have already stated above for the total level of cooperation, for higher degrees, the statistics are poorer, and the analysis does not lead to such clear-cut results.

Conclusions

To sum up, we have performed a large-scale experimental test of the hypothesis of network reciprocity (i.e., that the existence of a structure in the population may affect cooperation in social dilemmas). Our experiment shows that, when it comes to human behavior, the existence of an underlying network of contacts does not seem to have any influence in the global outcome. We want to stress that this conclusion applies only to human cooperation, and network reciprocity may still be relevant in other contexts

(e.g., microbiology) (25). Players seem to act by responding to the level of cooperation in their neighborhood, and this finding renders the network irrelevant. In addition, players behave in a moody manner, being significantly less likely to cooperate after a defection of their own. The levels of cooperation attained in a regular lattice and a highly heterogeneous network (hitherto thought to be a cooperation enhancer) are indistinguishable, and the responsive behavior of subjects seems to be independent of the number of neighbors that they have or the payoff differences that they observe. The results are in full agreement with the theoretical prediction in ref. 19; the fact that the key hypothesis in that model is that people behave in the way that we have just described provides additional support to our finding of moody conditional cooperation in networked PDs.

Our results have implications for policy-making when cooperation is a desired behavior. Although additional experiments with other social dilemmas still need to assess the range of applicability of our conclusions, the present study suggests that imposing a network structure might be a sterile effort. It should be noted, however, that this caveat does not imply that networking is useless to achieve cooperation—results would probably be very different if the network is allowed to be formed by the subjects as part of the game. Recent experiments on groups of up to 20 people (26, 27) strongly suggest this theory, but to the best of our knowledge, no large-scale experiments have been carried out to test this issue. However, the theoretical work in ref. 19 does not predict the slow decay of the cooperation level observed in the experiments, which we have conjectured arises from moody conditional cooperators becoming defectors in a generalized grim behavior. Such a change in the percentage of players using different strategies is not included in the theoretical model, and therefore, a next step would require accounting for such changes and if possible, justifying them within an evolutionary framework. Finally, given that, in our setup, players have to play the same action with all their neighbors, it is clear that our results should be related to Public Goods experiments. In fact, conditional cooperation was first observed in that context

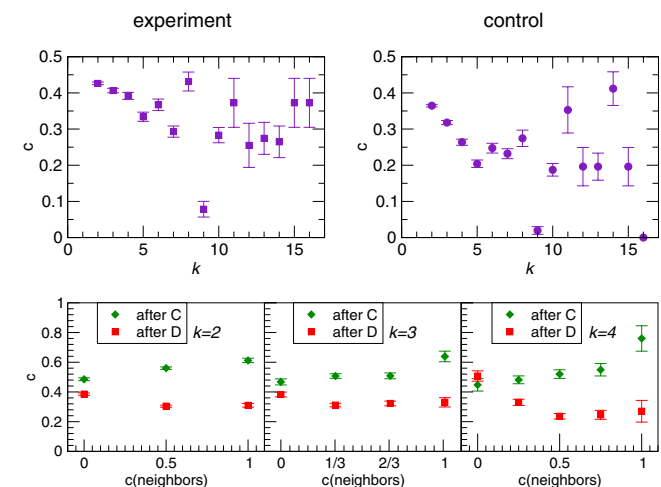


Fig. 5. Dependence of the strategies on the connectivity. *Upper* shows the cooperation level c as a function of the connectivity k_i in the heterogeneous network averaged over all rounds of the experiment (*Upper Left*) and the control (*Upper Right*) of the experiment. In *Lower*, we plot the frequency of cooperative actions of players with degree as indicated after they have cooperated or defected as a function of the fraction of cooperative actions in their neighborhood during the previous round; we also plot the experiment treatment in the heterogeneous network. Statistics are restricted to nodes of connectivity: $k = 2$ (*Lower Left*), $k = 3$ (*Lower Center*), and $k = 4$ (*Lower Right*).

(24). Our findings suggest that the moody version that we have found can also arise in Public Goods games. If that is the case, it is likely that network reciprocity does not apply to Public Goods games on networks. Hopefully, our experiment will encourage additional research in all these directions.

Materials and Methods

The experiment was carried out with 1,229 volunteers chosen among last year's high school students (17–18 y old) of 42 different high schools located throughout the geography of the Autonomous Region of Aragón, Spain. All of the students played through a web interface specifically created for the experiment (see Figs. S1–S3 and *SI Materials and Methods*) that was accessible through the computers available in the computer rooms of their respective schools. At least one teacher supervised the experiment in each computer room (which at most, had a maximum capacity of 20 students), preventing any interaction among the students. To further guarantee that potential interactions among students seating next to each other in the class did not influence the results of the experiment, the assignment of players to the different topologies was completely random. The colors used to code

the two available actions of the game were also selected randomly, also decreasing the likelihood that neighboring participants could influence each other. All participants went through a tutorial (included in *SI Materials and Methods*) on the screen, including questions to check their understanding of the game. When everybody had gone through the tutorial, the experiment began, lasting for approximately 1 h. The experiment assumed synchronous play; thus, we had to make sure that every round ended in a certain amount of time. This playing time was set to 20 s. At the end of the experiments, volunteers were presented a small questionnaire to fill in. Immediately after, all participants received their earnings and their show-up fee. Total earnings in the experiment ranged from 2.49 to 40.48 Euros.

ACKNOWLEDGMENTS. We thank E. Cauhé, J. J. Molinero, M. P. Pérez, and C. Viñas for their assistance during several phases of the project and J. Grujić for discussions during the software design phase. Work was supported by Fundación Ibercivis and Projects MOSAICO, PRODIEVO, FIS2008-01240, FIS2011-25167, FIS2009-13364-C02-01, FIS2009-12648-C03-02, and Complexity-NET RESINEE from Ministerio de Ciencia e Innovación (Spain); Project MODELICO-CM from Comunidad de Madrid (Spain); and a project to FENOL from Comunidad de Aragón (Spain).

- Fehr E, Fischbacher U (2003) The nature of human altruism. *Nature* 425:785–791.
- Pennisi E (2009) Origins. On the origin of cooperation. *Science* 325:1196–1199.
- Axelrod R (1984) *The Evolution of Cooperation* (Basic Books, New York).
- Rapoport A, Chammah AM (1965) *Prisoners Dilemma* (Univ of Michigan Press, Ann Arbor, MI).
- Axelrod R, Hamilton WD (1981) The evolution of cooperation. *Science* 211:1390–1396.
- Groves T, Ledyard J (1977) Optimal allocation of public goods: A solution to the free rider problem. *Econometrica* 45:783–809.
- Hardin G (1968) The tragedy of the commons. *Science* 162:1243–1248.
- Nowak MA (2006) Five rules for the evolution of cooperation. *Science* 314:1560–1563.
- Nowak MA, May RM (1992) Evolutionary games and spatial chaos. *Nature* 359:826–829.
- Szabó G, Fáth G (2007) Evolutionary games on graphs. *Phys Rep* 446:97–216.
- Roca CP, Cuesta JA, Sánchez A (2009) Evolutionary game theory: Temporal and spatial effects beyond replicator dynamics. *Phys Life Rev* 6:208–249.
- Santos FC, Pacheco JM (2005) Scale-free networks provide a unifying framework for the emergence of cooperation. *Phys Rev Lett* 95:098104.
- Gómez-Gardeñes J, Campillo M, Floría LM, Moreno Y (2007) Dynamical organization of cooperation in complex topologies. *Phys Rev Lett* 98:108103.
- Helbing D, Yu W (2010) The future of social experimenting. *Proc Natl Acad Sci USA* 107:5265–5266.
- Grujić J, Fosco C, Araujo L, Cuesta JA, Sánchez A (2010) Social experiments in the mesoscale: Humans playing a spatial prisoner's dilemma. *PLoS One* 5:e13749.
- Cassar A (2007) Coordination and cooperation in local, random and small world networks: Experimental evidence. *Games Econ Behav* 58:209–230.
- Kirchkamp O, Nagel R (2007) Naive learning and cooperation in network experiments. *Games Econ Behav* 58:269–292.
- Traulsen A, Semmann D, Sommerfeld RD, Krambeck HJ, Milinski M (2010) Human strategy updating in evolutionary games. *Proc Natl Acad Sci USA* 107:2962–2966.
- Gracia-Lázaro C, Cuesta JA, Sánchez A, Moreno Y (2012) Human behavior in Prisoner's Dilemma experiments suppresses network reciprocity. *Sci Rep* 2:325.
- Langer P, Nowak MA, Hauert C (2008) Spatial invasion of cooperation. *J Theor Biol* 250:634–641.
- Roca CP, Cuesta JA, Sánchez A (2009) The effect of population structure on the evolution of cooperation. *Phys Rev E* 80:46106.
- Ledyard JO (1995) *Public Goods: A Survey of Experimental Research*, *Handbook of Experimental Economics*, eds Nagel JH, Roth AE (Princeton Univ Press, Princeton), pp 111–251.
- Camerer CF (2003) *Behavioral Game Theory* (Princeton Univ Press, Princeton).
- Fischbacher U, Gächter S, Fehr E (2001) Are people conditionally cooperative? Evidence from a public goods experiment. *Econ Lett* 71:397–404.
- Velicer GJ (2003) Social strife in the microbial world. *Trends Microbiol* 11:330–337.
- Fehl K, van der Post DJ, Semmann D (2011) Co-evolution of behaviour and social network structure promotes human cooperation. *Ecol Lett* 14:546–551.
- Rand DG, Arbesman S, Christakis NA (2011) Dynamic social networks promote cooperation in experiments with humans. *Proc Natl Acad Sci USA* 108:19193–19198.

Supporting Information

Gracia-Lázaro et al. 10.1073/pnas.1206681109

SI Materials and Methods

1.1 Volunteer Recruitment and Treatment. The experiment was carried out with 1,229 volunteers chosen among last year's high school students (17–18 y old) of 42 different high schools located throughout the geography of the Autonomous Region of Aragón, Spain (capital is Zaragoza, the location of the University of Zaragoza); 34 high schools were in the province of Zaragoza, 5 high schools were in the province Huesca, and 3 high schools were in the province of Teruel. For the recruitment of the students, we contacted the coordinators of a program (Ciencia Viva, "Living Science") of the local government that supports the dissemination of science among public high schools of Aragón. Moreover, we also contacted many of the private schools of Zaragoza City to also offer to them the possibility of taking part in the experiment. In all cases, the experiment was referred to as a social experiment, and no one (including the high school teachers in charge of the coordination) knew in advance what the experiment was about (see below).

After the call for participation, we selected 1,300 volunteers. To satisfy ethical procedures, all personal data about the participants were anonymized and treated as confidential. Moreover, the Ethical Committee of the University of Zaragoza approved all procedures. On the day of the experiment, the aforementioned 1,229 volunteers showed up, with 541 males and 688 females representing 44.02% and 55.98% of the total number of players, respectively. Of the 1,229 students, 625 students played the game on a square lattice (274 males and 351 females to keep the male to female ratio), and 604 students played the game on an heterogeneous network. In the first topology, every player had $k = 4$ neighbors, whereas in the second topology, the connectivity varied between 2 and 16 using a distribution $\frac{N(k)}{N} = P(k) = Ak^{-2.7}$, with $A = (\sum_k P(k))^{-1}$.

All of the students played through a web interface specifically created for the experiment (see below) that was accessible through the computers available in the computer rooms of their respective schools. At least one teacher supervised the experiment in each computer room (which at most, had a maximum capacity of 20 students), preventing any interaction among the students. To further guarantee that potential interactions among students seated next to each other in the class do not influence the results of the experiment, the assignment of players to the different topologies was completely random. Hence, the odds of having two participants geographically close (i.e., of the same school and seating next to each other) who were also neighbors in the virtual topology was quite small. In addition, as described below, the colors used to code the two available actions of the game were also selected randomly, also decreasing the likelihood that neighboring participants could influence each other.

We describe the steps followed by each participant during the experiments. In short, all participants went through a tutorial on the screen, including questions to check their understanding of the game. When everyone had gone through the tutorial, the experiment began, lasting for approximately 1 h. At the end of the experiments, volunteers were presented a small questionnaire to fill in. Immediately after, all participants received their earnings and their show-up fee. Total earnings in the experiment ranged from 2.49 to 40.48 Euros.

1.2 Experimental Platform and Interface. The experiment was run using a fairly sophisticated web application specifically developed to this purpose. The application was made entirely using free software. It was developed in Ruby On Rails, a technology used by

other popular websites like Twitter, and has a MySQL database that stores all data needed to carry out the experiment and the subsequent analysis. MySQL is a freely available open-source relational database management system based on Structured Query Language, the most popular language for adding, accessing, and managing content in a database.

The application was designed to be used by three different user profiles. First, we have the players who were shown a detailed tutorial at the beginning (see below) for a better understanding of the interface and basis of the experiment. Second, there are teachers who were responsible for supervising students through their dedicated web monitors, facilitating the work of the central administrator, and contributing to the success of the experiment. Third, the administrators were responsible for controlling the game and everything that was happening in real time. The application, which was designed using technologies compatible with all platforms, was managed from a standard web browser. There was a last participant, a demon or process running in the background with the function to update the results and play instead of players who do not play within the specified time frame for each action.

Considering that the experiment required that around 1,300 students play online simultaneously, we used a server with enough power, and many optimizations were performed in terms of connections to the server, access to database, client-server data exchange, lightness of the interface, control logic, etc. The experiment started on December 20, 2011 at 10:00 CET. These steps were followed during the development of the experiment:

- i) Administrators opened the registration process.
- ii) Players (students) gradually registered.
- iii) After all students had registered, teachers informed the administrators through their screen.
- iv) As soon as the required number of participants had registered (this time took around 20 min), administrators blocked additional registrations and initiated the reading of the tutorial.
- v) Students and teachers read the tutorial.
- vi) Teachers informed (also through their screens) administrators that the reading was completed.
- vii) The experiment treatment began, which lasted 51 rounds.
- viii) Students played according to some predefined times (a maximum of 20 s per round to choose an action). During these steps, teachers controlled for any potential problem using a chat channel that connected them to the administrators. As mentioned above, if one student did not play within the 20 s given for each action, the demon played automatically (see below). The administrators were able to identify those students who was not playing and contact the teachers if the situation persisted. However, the experiment went smoothly, and no feedback to the professors for misbehavior was needed.
- ix) The experiment treatment finished, and a brief tutorial on the second experiment (control) was shown.
- x) After teachers and students had read the tutorial, the former notified the administrators.
- xi) Administrators started the control treatment, which lasted 59 rounds.
- xii) Students played as in the previous treatment.
- xiii) After the control treatment finished, volunteers were presented a short questionnaire to fill in.

xiv) All participants checked their earnings and were given their show-up fee.

1.3 Online Tutorial for Players. The following information is a translation of the Spanish original online tutorial (available on request). It is worth remarking that each player had a customized pair of colors and a corresponding number of neighbors. We refer to the latter as X (but X showed its actual value for each participant). As advanced above, to avoid framing effects, the two actions were always referred to in terms of colors (chosen randomly among a predefined set of possible pairs of colors), and the game was never referred to as Prisoner's Dilemma in the material handed to the volunteers. This information notwithstanding, subjects were properly informed of the consequences of choosing each action, and some examples were given to them in the introduction (see the tutorial text below). After every round, subjects were given the information of the actions taken by their neighbors and their corresponding payoffs. In all cases, the payoffs were properly normalized to avoid the possibility of guessing the number of connections of their neighbors. The instructions given here assume a given pair of colors (green and brown), but again, each participant saw the actual color assigned to him/her. Moreover, we took into consideration the possibility that some of the students were colorblind. In this sense, we provided clear instructions to avoid any possible error in the final results, specifying the order in which each color appeared on the screen and also using a combination of specifically selected colors (five different pairs); therefore, the probability of error was reduced to a minimum.

Page 1: This is an experiment designed to study how individuals make decisions.

You are not expected to behave in any particular way. Whatever you do will determine the amount of money you can earn. You have a written version of this direction, which you can check at any stage of the experiment. Please keep quiet during the experiment. If you need help, raise your hand and wait to be attended.

Page 2: Directions to participate in the experiment.

This experiment consists of TWO (2) parts. Each part consists of an undetermined number of ROUNDS (approximately between 50 and 70, but there might be more or less). Each part will last at most 35 min, but could finish before. In each part you will be able to earn different amounts of money, depending on the decisions that you and the rest of participants make in every round. The earning of each round is given in a monetary unit called ECU. When the experiment finishes, an exchange rate from ECUs to Euros will be established as a function of the number of participants. Your total earning in this experiment will be the accumulated earnings in all of the rounds of the two parts, plus a show-up fee.

Page 3: A round.

In each ROUND you will be placed in a node of a virtual NETWORK. In this network you will be linked to X (here, the actual number is shown to each participant) people, whom we shall refer to as "neighbors." Your neighbors will also be connected to other people. You will be one of those neighbors, but the rest of them will not necessarily be the same neighbors that you have.

You will never know who your neighbors are, and nobody will know if you are his/her neighbor either. The network is virtual. People around you in the room are not necessarily your neighbors.

Page 4: Decision to make in every round. Every round, each of the participants must choose a color: GREEN or BROWN. (As explained before, each participant sees the actual colors chosen for them. For clarity, we, henceforth, refer to green and brown,)

To choose a color you just have to click a button appearing in the screen.

Each time you choose a color (either blue or yellow) you will earn an amount of money which will depend on your and your X neighbors' choices.

If you choose GREEN and your neighbor also chooses GREEN, each receives 7 ECUs. If you choose GREEN and your neighbor chooses BROWN, you receive 0 ECUs and your neighbor 10 ECUs.

If you choose BROWN and your neighbor also chooses BROWN, each receives 0 ECUs.

If you choose BROWN and your neighbor chooses GREEN, you receive 10 ECUs and your neighbor 0 ECUs.

These rules are the same for all participants.

Page 5: Possible payoffs per neighbor.

In the following table each row corresponds to the decision you can make and each column correspond to one of your neighbors' decision.

Consider that:

you and each of your neighbors will globally earn more if you both choose GREEN (7 ECUs you/7 ECUs your neighbor); you will earn more if you choose BROWN and your neighbor chooses GREEN (10 ECUs you/0 ECUs your neighbor); but if both you and your neighbor choose BROWN you both will earn less (0 ECUs you/0 ECUs your neighbor) than if you both chose GREEN.

Page 6: This is the screen you will be seeing during the experiment (note that each participant actually sees the graph corresponding to his/her connectivity).

The central circle represents you, and the surrounding circles represent your virtual neighbors in that round.

On the right of the screen you will see two buttons: GREEN and BROWN.

Each round you must choose one of them clicking the corresponding button.

Page 7: These are some examples of what you could earn in a round.

Example 1: Imagine you choose GREEN, 3 of your neighbors choose GREEN and 1 chooses BROWN. In that round you will earn $3 \times 7 + 1 \times 0 = 21$ ECUs.

Example 2: In another round you choose BROWN, 2 of your neighbors choose GREEN and 2 choose BROWN. In that round you will earn $2 \times 10 + 2 \times 0 = 20$ ECUs.

Page 8: Round iteration.

Remember that each part will consist of an undetermined number of rounds.

Each round you will have up to 20 s to choose a color. After these 20 s, if you didn't choose, the system will choose for you. Whatever happens it will not affect the behavior of the system in the next rounds: you will be able to make your subsequent choices normally. (Don't worry: 20 s are more than enough to make a choice).

The round will not end until all participants have made their choice.

At the end of each round you will see a screen like this one. The central circle represents your choice (as given by the color) and your earning in this round. The surrounding circles represent your X neighbors' choices (represented by their colors) and their respective earnings in that round.

Your neighbors' earnings are given with respect to your number of neighbors. For example, you have 5 neighbors and one is Ferdinand (fictitious name). Ferdinand in turn has two neighbors: one is you and the other a stranger. If Ferdinand has won 10 ECUs in the last round, the gain of Ferdinand that you are shown is: $(10 \text{ ECUs}/2 \text{ neighbors of Ferdinand}) \times 5 \text{ neighbors of you} = 25 \text{ ECUs}$.

Note that what each of your neighbors has won depends on what you have chosen and also on what the neighbors of your neighbors have chosen.

Immediately after finishing a round there will be a new one, and then another one, and so on until you see a screen warning you about the end of that part of the experiment.

Page 9: Part I of the experiment.

In this part the system will randomly assign each participant to a given node of the virtual network.

This place will be kept fixed until this part ends.

This means that you will be interacting with **the same X neighbors** during all rounds of this part. Remember that in each round you must choose a color.

When this part finishes, you will be notified and will see the directions for the next part.

(Part I begins.)

Page 10: Part I of the experiment has finished.

Please keep quiet.

Part II will start in a few seconds.

Page 11: Part II of the experiment.

In this part, **before each round begins**, every participant will be moved to a **new** random node of the virtual network. Therefore, in general **you will likely have X new neighbors every round**.

This means **that the node you are in will be changing along the experiment**.

Thus, you will NOT be linked all rounds to the same X neighbors.

Page 12: The rules to make decisions every round are the same as in part I.

The only thing that is different is that your neighbors will most likely not be the same every round.

Remember:

Every round you have 20 s to make a choice.

The round finishes only when all participants have made their decisions.

At the end of each round you will be seeing a screen like in part I.

(Part II begins.)

Page 13: Part II of the experiment has finished.

Please, keep quiet.

The experiment has not finished yet.

You have to answer the following questionnaire.

Please, answer ALL questions in the questionnaire that you will be shown immediately.

(The questionnaire was shown and afterward they were notified how much they had earned and were to go to get paid.)

1.4 Synchronous Play and Automatic Actions. The experiment assumes synchronous play; thus, we had to make sure that every round ended in a certain amount of time. This playing time was set to 20 s, which was checked during the testing phase of the programs to be enough to make a decision, while at the same time, not too long to make the experiment boring to fast players. If a player did not choose an action within 20 s, the computer made the decision instead. This automatic decision was randomly chosen to be the player's previous action 90% of the times and the opposite action 10% of the times. We chose this protocol using previous testing performed in a similar experiment (1). Volunteers were informed that the computer would play for them if their decision took more than the prescribed timeout. However, they were not informed of the precise strategy used by the computer to avoid any bias in their own choices of strategy. In any case, for the reliability of the experiment, it is important that a huge majority of actions were actually played by humans and not by the computer. This quantity, when averaged over all rounds, yields that 90% of the actions were chosen by humans, regardless of the underlying network of contacts.

1.5 Questionnaires. At the end of the experiments, volunteers were presented a small questionnaire to fill in. The list of questions (translated into English) was as follows:

- i) Describe briefly how you made your decisions in part I (experiment).
- ii) Describe briefly how you made your decisions in part II (control).
- iii) Did you take into account your neighbors' actions?
- iv) Is something in the experiment familiar to you (yes/no)?
- v) If so, please point out of what it reminds you.
- vi) If you want to make any comment, please do so below.

The first three questions have a clear motivation, namely to see whether (possibly some) players did have a strategy to decide on their actions. Question 3 was intended to check whether players decided on their own or did look at their environment, because only in this last case do imitative or conditionally cooperative strategies make any sense. Questions 4 and 5 focused on the possibility that some of the players recognized the game as a Prisoner's Dilemma, because they had a prior knowledge of the basics of game theory. The final question just allowed them to enter any additional comment that they would like to make. We did not carry out a more detailed questionnaire to avoid the risk of many players' leaving it blank (the whole experiment was already very long).

SI Results and Discussion

Here, we present additional results aimed at supporting the findings shown in the text. As there, we will refer to the basic types of individuals found in the experiment as mostly cooperators (C; players who cooperate with a high probability regardless of the context), mostly defectors (D; players who defect with a high probability regardless of the context), and moody conditional cooperators (players whose action depends on their previous action as well as the level of cooperation in their neighborhood) (Fig. 3 A and B).

Fig. S4 shows the histograms of the number of players ranked according to the fraction of cooperative actions that they performed along the control phase in the lattice (Fig. S4, Left) and the heterogeneous network (Fig. S4, Right). The same results for the experimental phase can be found in Fig. 2 C and D. The comparison between the plots shows a large increase in the fraction of individuals that never or almost never cooperated in the control with respect to the experiment. This finding is likely to be a consequence of the fact that, in the experiment, there is an initial amount of cooperation well above 50%, which is not

the case in the control. At the other extreme of the plots, the (small) amount of highly cooperative players remains approximately the same, indicating that their motivation has nothing to do with having or not having a fixed environment for their interactions. The general picture thus arising from the control part is that there is not much cooperation, and the majority of players do not cooperate other than occasionally.

However, Fig. S5 displays the time evolution of the distribution of cooperative actions in the experimental part. The histograms show the players' frequency as a function of the fraction of cooperative actions along successive 10-round periods corresponding to the experimental phase in the lattice (Fig. S5, *Left*) and the heterogeneous network (Fig. S5, *Right*). The results show evidence of some degree of learning as the experiment progresses. Indeed, the number of people who cooperate never or rarely increases with time. This finding would be consistent with the decay of cooperation shown in Fig. 24; although the first quick drop in cooperation would be explicable within a computer model with a fixed proportion of D, C, and moody conditional cooperators, the second part of the evolution, a much slower decay, is inconsistent with such a model and must then come from changes in the proportion of the different types of players.

The phenomenon that we have just described can also be shown in a different manner, namely by monitoring the evolution of mostly defectors during both the experimental and control parts of the experiment. Fig. S6 represents the fraction of agents with probability to cooperate that is below a given threshold (indicated

on the right) at every round (time t). To calculate this quantity, we have taken into account the actions of the players during the previous 10 rounds. The results obtained show an increasing trend (more evident for the experimental phase) (Fig. S6, *Upper*) for both the square lattice and the heterogeneous network, which confirms the tendency of the players to learn that they should defect as time goes on.

We also report on the statistical analysis that we carried out about the experimental data. To determine whether the likelihood to cooperate differs significantly in the two studied networks, we use the Kolmogorov–Smirnov test for the two datasets. We take, as a first sample, the distribution of the probability to cooperate in the lattice cumulated over all rounds of the experimental phase. The second sample used as input for the Kolmogorov–Smirnov test corresponds to the same distribution but for the heterogeneous network. These distributions are represented in Fig. 2. The maximum difference between the cumulative distributions for the experimental phase is 0.1071, with a corresponding value for $P_{KS} = 0.995$. The statistics of both samples, together with the ones corresponding to the control phase in Fig. S4, are summarized in Table S1.

Finally, Table S2 summarizes the statistical fits (obtained from a weighted least squares regression) of the conditional probability P to cooperate conditioned on the player's action in the previous round ($X =$ after C or after D) and the density ρ of cooperators in the players' neighborhoods during the previous round. Fits are defined by $P(C|X, \rho) = a + b\rho$. The data fitted correspond to the results shown in Fig. 3 *A* and *B*.

1. Grujić J, Fosco C, Araujo L, Cuesta JA, Sánchez A (2010) Social experiments in the mesoscale: Humans playing a spatial prisoner's dilemma. *PLoS One* 5:e13749.

GANANCIAS POSIBLES POR CADA VECINO:

En la siguiente tabla, cada fila corresponde a la elección que podría hacer usted y cada columna, a la elección que podría hacer uno de sus vecinos.

	COOPERACIÓN (7 ECU)	DEFECCIÓN (0 ECU)
SU ELECCIÓN (7 ECU)	7	0
SU ELECCIÓN (0 ECU)	10	0

Por lo tanto, usted y cada vecino suyo ganan en conjunto más si ambos eligen ■ (7 ECU usted / 7 ECU su vecino); usted gana más si elige ■ y su vecino elige ■ (10 ECU usted / 0 ECU su vecino); usted no gana nada si elige ■ y su vecino elige ■ (0 ECU usted / 10 ECU su vecino); pero si usted y su vecino eligen ■ ganan menos (0 ECU usted / 0 ECU su vecino) que si ambos hubieran elegido ■.

ANTERIOR
SIGUIENTE

Fig. S1. Snapshot of the experimental software.

Esta será la pantalla que verá durante el experimento:

Ganado en Ronda Anterior



Ronda

Para jugar esta ronda seleccione un color

El círculo central es usted y los círculos a su alrededor, sus 15 vecinos virtuales en esa ronda.

En la derecha de la pantalla aparecen dos botones, el ■ y el ■.

En cada ronda debe elegir color pinchando el botón que desee.

ANTERIOR

SIGUIENTE

Fig. S2. Snapshot of the experimental software. Note that the payoffs shown do not correspond to any real situation but simply illustrate how they were seen by the subjects.

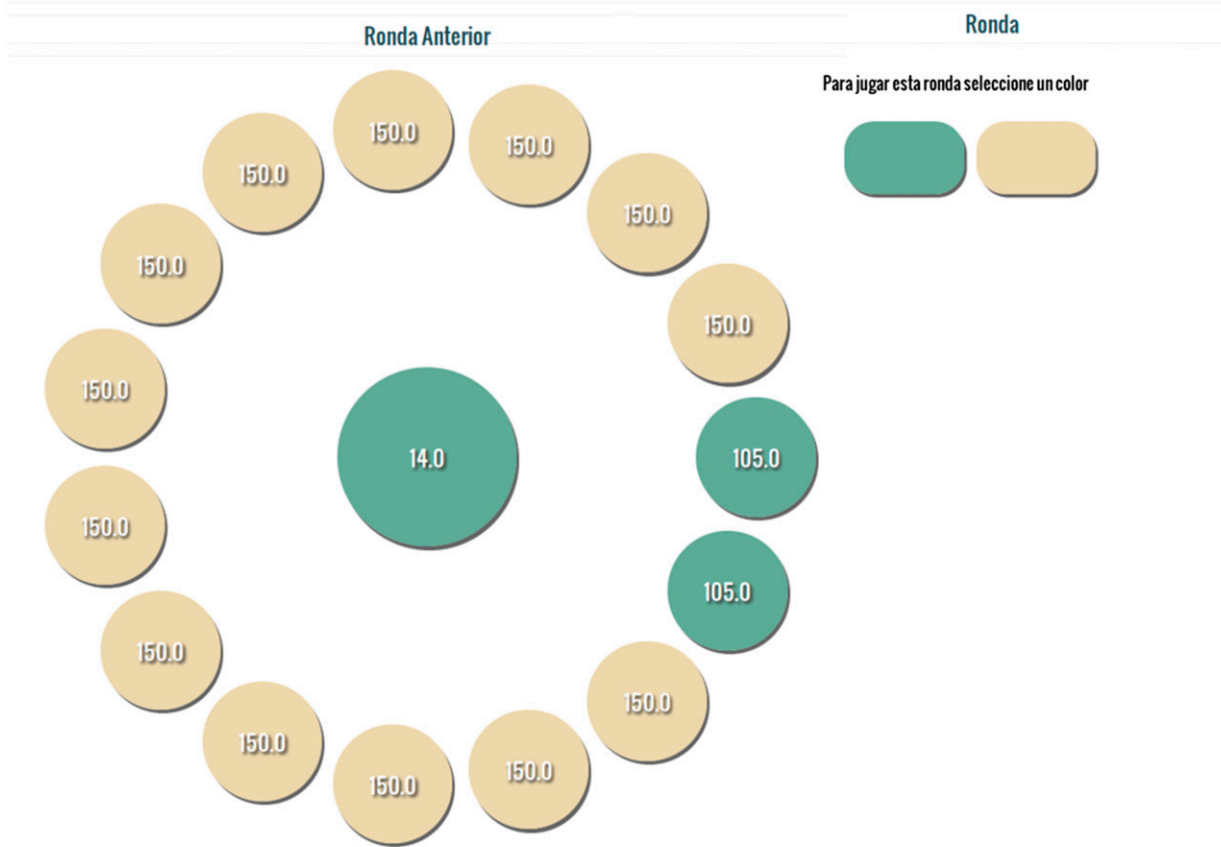
REPETICIÓN DE RONDAS

Recuerde que en cada parte habrá un número indeterminado de rondas.

En cada ronda usted tiene hasta 20 segundos para elegir color. Pasados los 20 segundos, el sistema elegirá por usted, aunque después usted podrá seguir eligiendo sin problemas en las rondas siguientes. (No se preocupe, 20 segundos deberían sobrarle para elegir).

La ronda no termina hasta que hayan elegido todos los participantes.

Al finalizar la ronda aparecerá una pantalla como esta:



En el círculo central está usted, con el color, la ganancia que ha obtenido en esta ronda y la ganancia total. En los círculos a su alrededor, aparecen sus 15 vecinos en esa ronda, con el color que cada uno ha elegido y la cantidad que ha ganado por vecino por el número de vecinos que usted tiene.

Por ejemplo, usted tiene 15 vecinos y uno de sus vecinos es Ferdinand (nombre ficticio), y Ferdinand tiene a su vez a dos vecinos: uno es usted y otro un desconocido. Si Ferdinand ha ganado 10 ECUs en la última ronda, la ganancia de Ferdinand que a usted se le muestra es

$(10 \text{ ECUS} / 2 \text{ VECINOS DE FERDINAND}) * 15 \text{ VECINOS SUYOS} = 75 \text{ ECUs}$

Note que lo que cada vecino suyo ha ganado depende de lo que usted ha elegido y de lo que han elegido los restantes vecinos de su vecino.

Inmediatamente después de terminar una ronda, habrá otra ronda, y después de ésta, otra más y así sucesivamente hasta que reciba el aviso por pantalla que la parte del experimento ha terminado.

ANTERIOR

SIGUIENTE

Fig. S3. Snapshot of the experimental software.

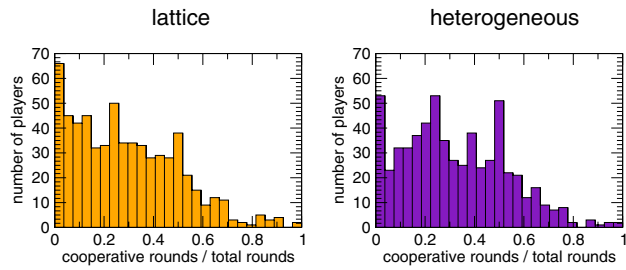


Fig. 54. Distribution of cooperative actions in the control. We represent the number of players that cooperated during the given number of rounds (normalized by the total number of rounds played). The results correspond to the control phase. Similar results were presented in Fig. 2.

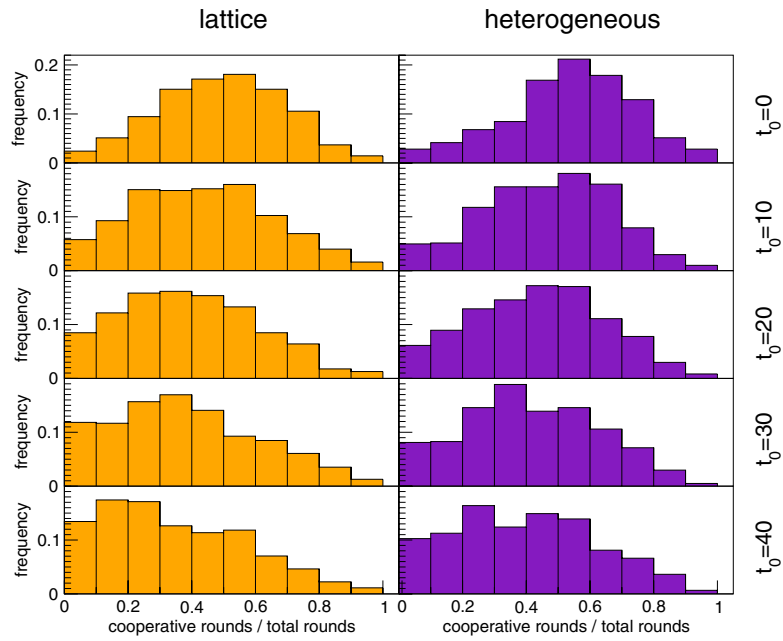


Fig. 55. Time evolution of the distribution of cooperative actions. The different panels show how frequently players cooperated in different time periods. The results correspond to the first treatment (experiment). Rows represent periods 1–10 ($t_0 = 0$), 11–20 ($t_0 = 10$), 21–30 ($t_0 = 20$), 31–40 ($t_0 = 30$), and 41–50 ($t_0 = 40$) as indicated.

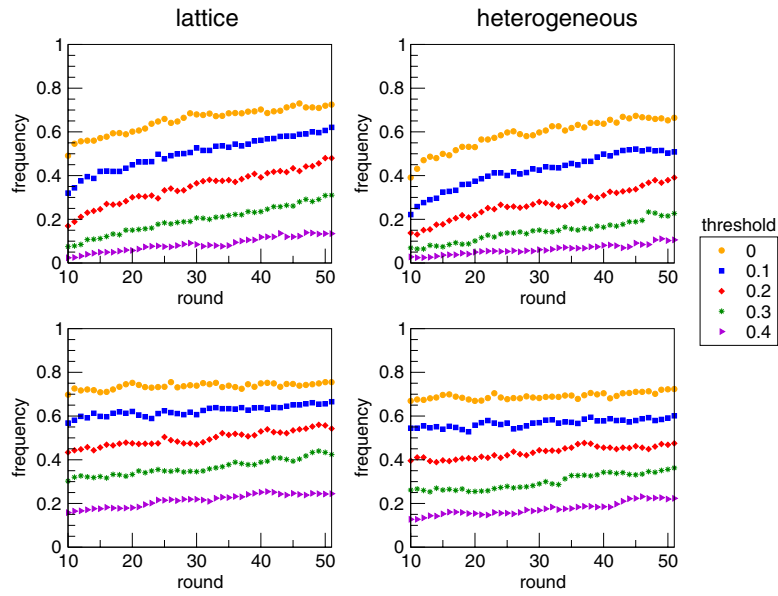


Fig. S6. Evolution of the fraction of mostly defectors. Fraction of agents with a cooperation probability lower than a given threshold as a function of t (=round) according to their cooperative actions through the previous 10 rounds for different values of the $threshold = 0, 0.1, 0.2, 0.3, 0.4$. Results for the (Left) lattice and (Right) heterogeneous network. Two treatments: (Upper) experiment and (Lower) control.

Table S1. Statistics of the distribution of the probability to cooperate cumulated over all rounds of the experimental and control phases in both networks (additional details in the text)

	Experiment		Control	
	Lattice	Heterogeneous	Lattice	Heterogeneous
Mean	0.03703	0.03703	0.03226	0.03226
95% confidence interval	(0.02434–0.04974)	(0.02335–0.05072)	(0.02549–0.04858)	(0.02607–0.04800)
SD	0.03210	0.03459	0.02918	0.02772
High	0.0976	0.104	0.106	0.0878
Low	0	0	0	0
Third quartile	0.06560	0.06126	0.05440	0.05795
First quartile	0.006400	0.006623	0.006400	0.01159
Median	0.04000	0.03146	0.0448	0.03808
Median absolute deviation	0.02844	0.02937	0.02495	0.02275

Table S2. Values of the fitting parameters for the results shown in Fig. 3 A and B

	Lattice		Heterogeneous	
	Fig. 3A	Fig. 3B	Fig. 3A	Fig. 3B
After C	0.457 ± 0.015	0.122 ± 0.034	0.475 ± 0.016	0.126 ± 0.039
After D	0.350 ± 0.021	-0.149 ± 0.050	0.309 ± 0.069	-0.0269 ± 0.035

Fits are defined by $P(C|X, \rho) = a + b\rho$, with $X = \text{after C or after D}$. Additional details in the text.