

Fusing enacted and expected mimicry generates a winning strategy that promotes the evolution of cooperation

Ilan Fischer^{a,1}, Alex Frid^b, Sebastian J. Goerg^{c,d}, Simon A. Levin^{e,1}, Daniel I. Rubenstein^e, and Reinhard Selten^f

^aDepartment of Psychology and ^bEdmond J. Safra Brain Research Center, University of Haifa, Haifa 31905, Israel; ^cDepartment of Economics, Florida State University, Tallahassee, FL 32306; ^dMax Planck Institute for Research on Collective Goods, D-53113 Bonn, Germany; ^eDepartment of Ecology and Evolutionary Biology, Princeton University, Princeton, NJ 08544; and ^fDepartment of Economics, University of Bonn, 53113 Bonn, Germany

Contributed by Simon A. Levin, May 8, 2013 (sent for review December 11, 2012)

Although cooperation and trust are essential features for the development of prosperous populations, they also put cooperating individuals at risk for exploitation and abuse. Empirical and theoretical evidence suggests that the solution to the problem resides in the practice of mimicry and imitation, the expectation of opponent's mimicry and the reliance on similarity indices. Here we fuse the principles of enacted and expected mimicry and condition their application on two similarity indices to produce a model of mimicry and relative similarity. Testing the model in computer simulations of behavioral niches, populated with agents that enact various strategies and learning algorithms, shows how mimicry and relative similarity outperforms all the opponent strategies it was tested against, pushes noncooperative opponents toward extinction, and promotes the development of cooperative populations. The proposed model sheds light on the evolution of cooperation and provides a blueprint for intentional induction of cooperation within and among populations. It is suggested that reducing conflict intensities among human populations necessitates (i) instigation of social initiatives that increase the perception of similarity among opponents and (ii) efficient lowering of the similarity threshold of the interaction, the minimal level of similarity that makes cooperation advisable.

conflict resolution | prisoner's dilemma | SERS | TFT | WSLS

Cooperation among individuals is a key determinant of prosperous populations (1, 2), yet cooperative individuals are prone to exploitation by competitive and hostile opponents. How then can organisms evolve cooperative behaviors? Theoretical and empirical evidence suggests that a key to this riddle may reside in the reliance on mimicry, similarity, and imitation. Coevolution of various species demonstrates how mimicry of color, movement, and shape provides powerful protection against predators, helps in luring prey (3), and enables parasites to attain cooperative behavior from their hosts (4). Among humans, similarity induces interpersonal attraction (5), affiliation (6), rapport (7), generosity (8), and cooperation (9–11). Imitation provides the means for learning and coping with social and strategic structures (12–17), whereas homophilous imitation, the imitation of others who are similar to one's self, generates stable and homogeneous populations (18). Simulation studies of the evolution of cooperation using the prisoner's dilemma (PD) paradigm (Table 1) revealed two powerful strategies, each representing a different aspect of mimicry and imitation. Tit for tat (TFT) (19), which looks backward and mimics opponent's previous move; and win–stay, lose–shift (WSLS) (20), which looks ahead by initiating a cooperative move whenever both opponents have made identical choices. Animating both strategies, we may describe TFT as enacting mimicry and WSLS as expecting similar opponents to make matching cooperative choices. Nevertheless, not any instance of observed similarity indicates the existence of a genuine and sufficiently high level of similarity. The occurrence of similar choices may result from stochastic processes, be strategically manipulated, or result from a rather low level of similarity. Subjective expected relative similarity (SERS)

theory (10) shows that every PD game, as well as other similarity sensitive interactions (11), has a similarity threshold that distinguishes among insufficiently and sufficiently high levels of an opponent's similarity (Table 1). As shown in the present work, fusing the principles embedded in TFT and WSLS while conditioning their function on SERS generates an overwhelmingly powerful and cooperative strategy. The merged mechanism is backward- and forward-looking; it fights noncooperative opponents and promotes the evolution of cooperation in its environment.

Mimicry and Relative Similarity

Combining the principles embedded within TFT, WSLS, and SERS, we derive a strategy of mimicry and relative similarity (MaRS) that is applicable to repeated PD conflicts as well as to other strategic interactions that are sensitive to an opponent's similarity (11). MaRS is a simple strategy that switches among three states of mimicry—enacted, expected, and excluded—allowing it to select an optimal response for every opponent and stage of the game (Table 2). Enacted mimicry resembles TFT, expected mimicry resembles WSLS's shift toward cooperation, and excluded mimicry turns off all cooperative tendencies and transforms MaRS into a defector. The transitions from one state of mimicry to another are guided by the current extent of similarity between the opponent and itself. MaRS monitors two types of similarity by updating two respective registries in a first-in-first-out manner: (i) the passive similarity registry that reflects the proportion of similar choices, regardless of whether they were cooperative or hostile; and (ii) the reactive similarity registry that reflects the opponent's propensity to reciprocate switches toward cooperation that were initiated by MaRS. The former is updated by using values of 1 and 0 for similar and dissimilar choices, respectively. The latter is updated by values of 1 and 0 for reciprocated and nonreciprocated switches, respectively. Nevertheless, to assure further the genuine nature of reciprocated cooperative switches, MaRS continues to monitor the succeeding moves. If the opponent switches back to defection, the reactive registry is updated with -1 , thus canceling out the 1 earned for the previous reciprocation. Continuously monitoring both similarity types allows computing two indices and applying them as criteria for the selection and transition among the states of mimicry. As derived from SERS, the indices are weighed in respect to the similarity threshold of the game, p_s^* . This threshold is the precise level of similarity that accounts for an identical expected payoff, regardless whether one

Author contributions: I.F. designed research; I.F. performed research; A.F. developed computer algorithms; S.J.G. and R.S. contributed learning algorithms and game theoretic perspectives; S.A.L., and D.I.R. contributed ecological and evolutionary perspectives; and I.F. wrote the paper.

The authors declare no conflict of interest.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. E-mail: ifischer@psy.haifa.ac.il or slevin@princeton.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1308221110/-DCSupplemental.

Table 1. PD game and SERS solution

Player 1	Player 2	
	Cooperate	Defect
Cooperate	R, R	S, T
Defect	T, S	P, P

Payoffs for player 1(left value in each cell) and player 2 (right value in each cell) are denoted as T, R, P, and S, where $T > R > P > S$ and $2R > S + T$, to define the PD structure. Introducing p_s , the probability that indicates the extent of similarity with the opponent, which can be estimated from the relative frequency of both players choosing the same action (cooperate/cooperate or defect/defect) in previous encounters of a repeated game, we denote the expected payoffs for cooperation ($p_s \times R + (1 - p_s) \times S$) and defection ($p_s \times P + (1 - p_s) \times T$), respectively. Hence similarity-focused and payoff-maximizing individuals may take the risk associated with cooperation whenever the similarity with the opponent exceeds the threshold: $p_s^* = (T - S)/(T - S + R - P)$. This threshold minimizes the risks associated with the choice of cooperation and is hence used by MaRS agents for the evaluation of passive and reactive similarity indices.

chooses to cooperate or to defect [referring to the PD game payoffs of T,R,P, and S, where $T > R > P > S$ and $2R > S + T$, $p_s^* = (T - S)/(T - S + R - P)$; Table 1]. If similarity with the opponent exceeds this threshold, cooperation provides a higher expected value than defection. If similarity fails to reach the threshold, the choice of defection promises a better expected value. As the choice among the three mimicry modes also involves taking the risk associated with cooperation (e.g., mimicking a cooperative move in enacted mimicry or initiating cooperation while using expected mimicry), p_s^* provides an adequate criterion for assessing the magnitude of both similarity indices.

The operation of MaRS can be described in terms of its switches among the three mimicry modes (Fig. 1) or in terms of its transitions among the choices of cooperation and defection (Fig. 2). *SI Appendix, Tables S1–S3*, provide further details by showing a step-by-step example of a repeated PD game between MaRS and a random playing opponent, as well as a comparison with TFT and WSLY playing against the same sequence of random moves. *SI Appendix, Appendix 1*, describes the function of MaRS in a pseudoprogramming code. Note that the capacity of MaRS’s similarity registries is a free parameter. Large capacities reflect long memory spans and slow adaptation rates, whereas small capacities reflect short memory spans and fast adaptation rates. Although it is not easy to predict which memory span will provide an optimal performance when interacting with unfamiliar opponents, delaying the use of the reactive memory for a few interactions helps to

Table 2. MaRS strategy

	Last move in period <i>n</i>		Similarity indices		Next moves by MaRS	
	Opponent	MaRS	Passive	Reactive	<i>n</i> + 1	<i>n</i> + 2
1. First interaction	NA		NA		●	
2. Enacted mimicry						
a	■	●	+	–	●	
b	■	●	+	+/-/?	●	
c	■	○	+	+/-/?	●	
3. Expected mimicry	■	●	+	+/?	●	●
4. Excluded mimicry	□	○	–	+/-/?	●	

The MaRS strategy comprises enacted, expected, and excluded mimicry. The moves of MaRS and its opponent are denoted by circles and squares. Cooperation and defection are indicated by blue and red colors, respectively. Empty circles or squares indicate that the actual choice is not important. “+”, “–”, and “?” indicate that the current similarity index (passive or reactive) is above or below the similarity threshold p_s^* of the game, or that its value is yet unknown, where $p_s^* = (T - S)/(T - S + R - P)$; “*n*”, “*n* + 1”, and “*n* + 2” denote subsequent moves. Note that only expected mimicry defines the next two moves. NA, not available.

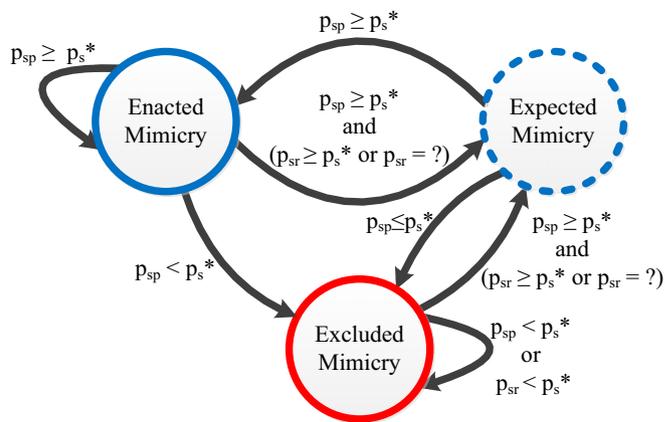


Fig. 1. The transition of MaRS among three mimicry modes: enacted, expected, and excluded. Each transition is conditioned upon the detection of up-to-date similarity levels by the passive similarity index p_{sp} , the reactive similarity index p_{sr} , or both. The criterion for sufficient similarity is the similarity threshold p_s^* , which is derived from the game’s payoff matrix, whereby $p_s^* = (T - S)/(T - S + R - P)$. “ $p_{sr} = ?$ ” indicates that the extent of reactive similarity has not yet been determined. Note that the transitions start at the third interaction, at which point a mimicry type has already been determined and MaRS considers its next move.

prevent premature conclusions in regard to the reactive similarity with the opponent. Such a delay reflects the need of organisms to orient themselves and adapt to the behavior of their opponents. The simulations presented hereafter approximate the capacities typically attributed to human working memory, which holds 7 ± 2 items (21) and delays the use of the reactive memory until it contains five items.

Testing MaRS in Simulated Behavioral Niches

To test the performance of MaRS, we constructed several artificial behavioral niches composed of simulated strategies and learning algorithms that interact in a repeated PD roundtable competition over many generations. At the termination of each generation, the pooled gains of each strategy are transformed into a new proportion of agents that enact the same strategy during the next generation. This procedure simulates selection pressures, generating a growth in the number of successful agents and a decrease in the number of less successful ones.

We populate the simulated niches with two sets of strategies. The first set consist of strategies, typically associated with the PD game, comprising nonreactive cooperators (all C), nonreactive defectors (all D), random playing agents (random), TFT (17), and WSLY (18).

The second set consists of five prominent stimulus- and belief-based learning algorithms: (i) reinforcement learning (22, 23), a simple stimulus-learning mechanism that chooses each action with a probability proportional to the relative gains earned while using it in previous encounters; (ii) payoff-sampling learning (24, 25), which randomly samples payoffs obtained following the use of each choice; (iii) impulse-balance learning (25, 26), which examines the gap between obtained and forgone payoffs and generates impulses toward better choices; (iv) fictitious play (27, 28), which maximizes the payoffs against the entire history of opponent’s previous choices; and (v) action-sampling learning (25), which optimizes its actions against a random sample of the opponent’s previous choices and thus constitutes a stochastic version of fictitious play.

Fig. 3 depicts the evolutionary course developing from six different initial population compositions. Fig. 3A depicts the evolutionary course of an initial population that comprises of six strategies: MaRS and five other strategies typically associated with the PD game. After the extinction of the nonreactive defectors and

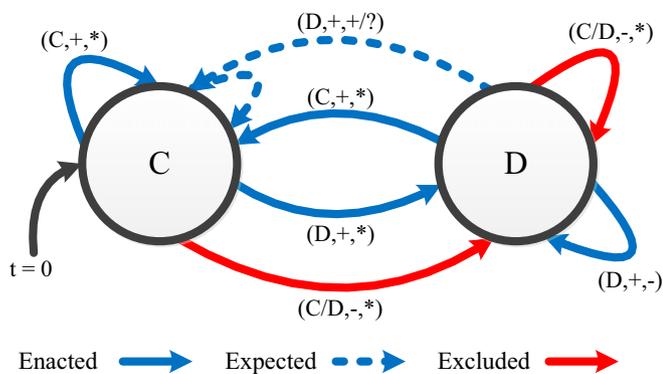


Fig. 2. The transitions of MaRS among the choices of cooperation (marked as "C") and defection (marked as "D") for a repeated PD game. The arrows are marked with ordered triples (opponent's last move, passive similarity rank, reactive similarity rank) defining the conditions for each type of transition. Plus sign, minus sign, and question mark indicate whether the current value of the similarity index, either p_{sp} or p_{sr} , is above (or identical) or below the similarity threshold p_s^* , or whether its value has not yet been determined, whereby $p_s^* = (T - S)/(T - S + R - P)$. An asterisk accounts for all possible states of the similarity indices; a question mark indicates that the value of the index is yet undefined. Each triplet is labeled with the relevant type of mimicry: enacted, expected, or excluded. Note that, for all cases of enacted mimicry, the passive similarity value is above the similarity threshold (+), and that, for all cases of excluded mimicry, the passive similarity value is below the similarity threshold (-). Also note that, in the case of expected mimicry, MaRS cooperates twice. $t = 0$ indicates that the game is about to begin.

the random playing agents, the proportion of all cooperative strategies stabilizes to form a cooperative population in which most agents are MaRS, followed by TFT, WSLs, and nonreactive cooperators. Fig. 3B depicts the evolutionary course of an identical set of six strategies, but with an enlarged proportion of unconditional defectors and random agents. Although, initially, the unconditional defectors prosper, soon their course of development changes. After the extinction of the nonreactive defectors and the random playing agents, the proportion of all remaining strategies stabilizes to form a cooperative population with a meaningful majority of MaRS agents, followed by TFT, a small proportion of WSLs, and a tiny proportion of nonreactive cooperators (i.e., all C). Note that, although Fig. 3A and B depicts a convergence to two different equilibria, they point to the existence of a continuum of such states. To outline their range consider the existence of any proportion of the four cooperative strategies (MaRS, TFT, WSLs, and all C) while setting the number of all D and random playing agents to zero. As all cooperative strategies cooperate in each and every encounter with all others, no evolutionary dynamics develop and the proportions of the cooperative strategies in the initial population do not change. However, modifying the original population by adding all D and random agents is expected to diminish the fitness of the cooperative strategies and account for changes in the composition of the population. A crucial factor in this dynamics is the function of MaRS, which identifies the intruders, acts to diminish their number (SI Appendix, Fig. S2), and restores the cooperative nature of the population. When the hostile and random agents have been removed, the population converges, as depicted in Fig. 3A and B, and its exact composition depends on the proportions of the initial population and the number of generations that passed before the convergence has been reached.

Fig. 3C depicts the evolutionary course of an initial population with equal proportions of MaRS and five learning strategies. After a few generations, all learning strategies go extinct and MaRS agents take over the entire population. Fig. 3D depicts the evolutionary course of an extended and rather hostile population that also contains many nonreactive defectors and random playing

agents. As before all strategies, including all D and random, go extinct and the population becomes fully cooperative, populated exclusively by MaRS agents. Fig. 3E depicts the evolutionary course of an initial population with equal proportions of all 11 strategies. Most agents go extinct within 10 generations, and the only surviving strategies are MaRS, followed by TFT and a rather small proportion of unconditional cooperators. The converged population has become completely cooperative. Fig. 3F depicts the evolutionary course initially composed of the same 11 types of agents, but with an enlarged proportion of unconditional defectors and random behaving agents. Although, initially, the unconditional defectors prosper, soon their course of development changes and the cooperative agents take over. Finally, all agents go extinct except for three cooperative strategies: MaRS followed by TFT and a small proportion of WSLs agents. As before, the converged population has become completely cooperative.

Although the aforementioned simulations reveal the performance of MaRS while confronting many other prominent rivals, embedded in various environments, all simulations were conducted while using a single matrix of the PD game, whereby $T = 5$, $R = 3$, $P = 1$, and $S = 0$, reflecting a similarity threshold of $p_s^* = 0.714$. Nevertheless, there exists a direct relation between higher thresholds and MaRS's advantage over other strategies (SI Appendix, Fig. S1). The reason for the increased advantage of MaRS stems from the costs it pays for misidentifying an opponent for being sufficiently similar, hence applying enacted mimicry instead of excluded mimicry, or for attempting to assess an opponent's reactive similarity by applying expected mimicry. When the payoff matrix has a low similarity threshold, short and hence nonrepresentative samples of random opponents or short samples of strategies that are almost sufficiently similar may seem to reveal sufficient similarity and motivate risky moves. Such nonrepresentative samples become scarcer when the required similarity threshold has higher values. Importantly, this does not lead to the conclusion that MaRS performs poorly when encountering random opponents. In fact, a key advantage of MaRS is that it is not easily misled by random playing agents. If the passive similarity with the opponent is insufficiently high, sporadic cooperative moves will not be reciprocated, nor will they motivate MaRS to switch to expected mimicry. Even if a random opponent has exhibited several instances of sufficient similarity, motivating MaRS to apply expected mimicry, the consequences will be monitored by the reactive similarity registry and are likely to prevent such moves in future encounters (SI Appendix, Table S1). In fact, adding random agents to a heterogeneous population actually increases the advantage of MaRS (SI Appendix, Fig. S2).

Discussion

Fusing enacted and expected mimicry generates a powerful and cooperative mechanism that enhances fitness and reduces the risks associated with trust and cooperation. When conflicts take the form of repeated PD games, individuals get the opportunity to learn and monitor the extent of similarity with their opponents. They can then react by choosing whether to enact, expect, or exclude mimicry. This rather simple behavior has the capacity to protect individuals from exploitation and drive the evolution of cooperation within entire populations. Like TFT, MaRS paves the way for the induction of cooperation and supports the survival of other cooperative strategies (29). The existence of MaRS in heterogeneous populations helps those cooperative strategies that do not have the capacity of MaRS to combat hostile and random opponents. Despite the fact that MaRS cannot prevail in a duel with an unconditional defector, interacting within heterogeneous populations allows MaRS to fight unpredictable and hostile strategies and cooperate with cooperative ones, including itself. This capacity enables a handful of MaRS agents to invade a population of unconditional defectors (e.g., adding 2% of MaRS agents, with identical parameters to those shown in Fig. 3, to a population of 1,000 unconditional defectors results in a

complete extinction of the defectors after 40 generations). The operation of MaRS promotes cooperation, minimizes the extent of exploitation, and accounts for high fitness levels. Even if the algorithm is exposed, MaRS agents do not face additional risks; their opponents can fake neither passive or reactive similarity, nor can they escape the reciprocity of enacted mimicry. Having the capacity to learn and adapt to opponents' strategies, MaRS can be thought of as being endowed with a "theory of mind" (30). It distinguishes among various types of opponents and games and may hence respond differently to identical moves conducted by different opponents. This adaptive response prevents other strategies from enforcing a claim for an unfair share of the payoffs (30). Nonetheless, being dependent on its memory to

characterize an opponent's similarity implies having greater vulnerability to errors. When an error has occurred, it is carried forward through memory and may influence several future moves of MaRS. If encountering other memory-dependent strategies, the errors are likely to have an impact also on their future conduct and add to the complexity of the developing interactions.

Although a comprehensive study of various error types and their impacts on heterogeneous strategy populations is a worthwhile task, it is beyond the scope of the present manuscript. Instead, we make three comments in respect of genuine populations:

First, MaRS as materialized in the present study begins each repeated interaction with empty memory registries and a slight cooperative bias; however, biological and social environments

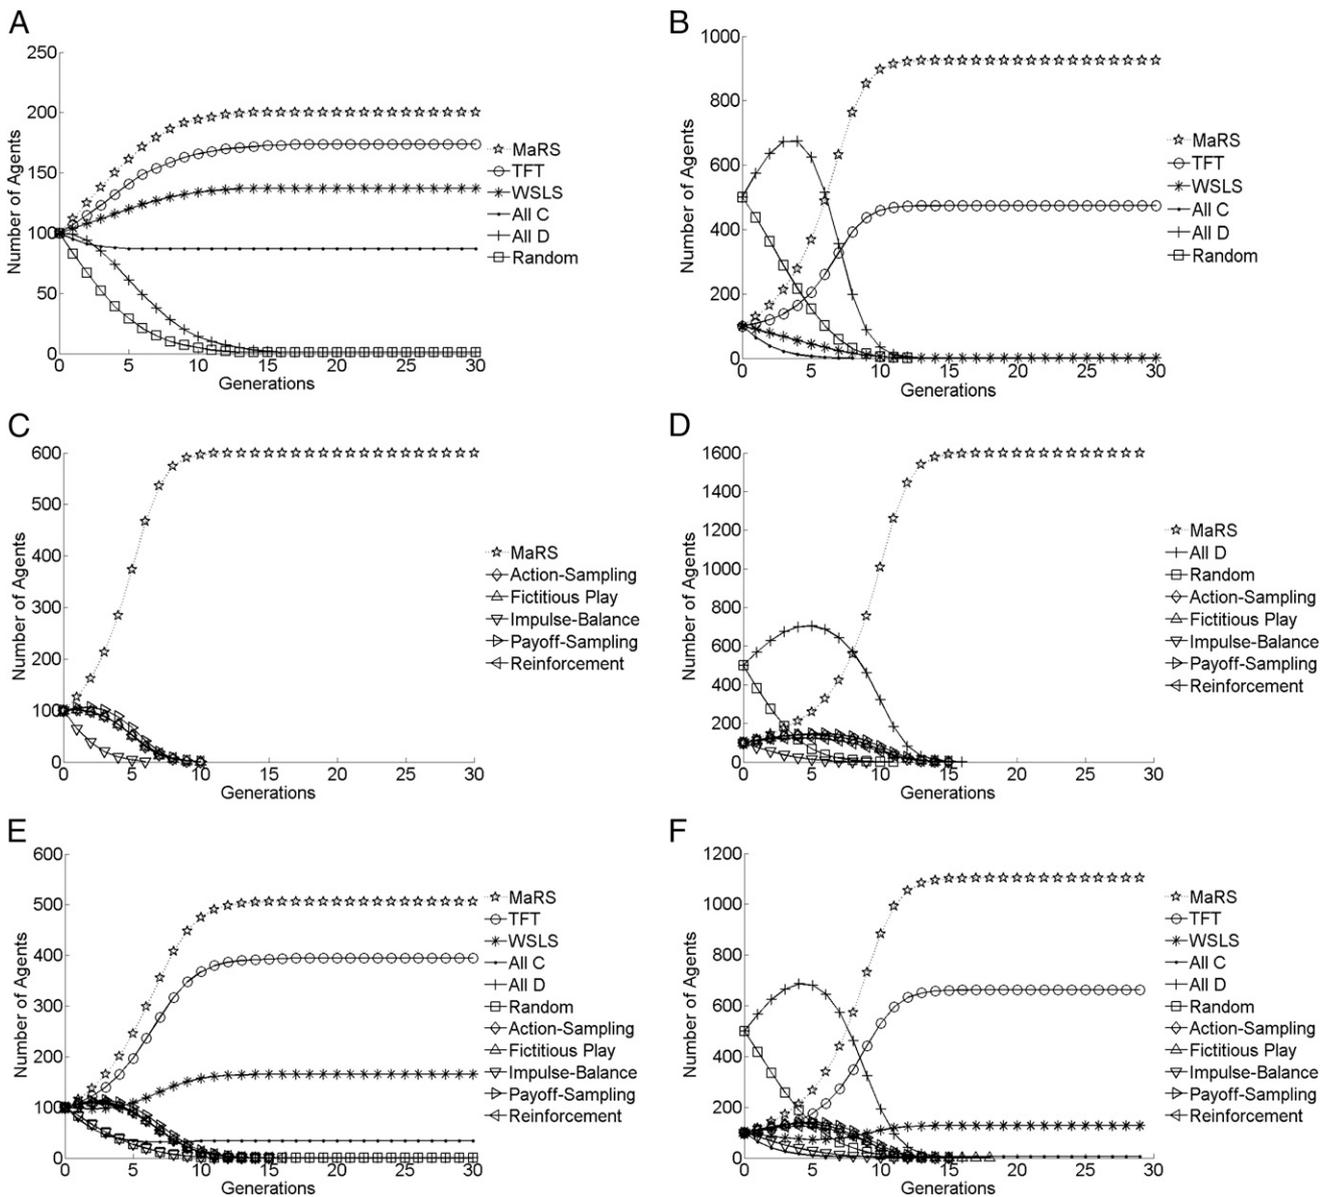


Fig. 3. Evolutionary simulations of six behavioral niches initially populated with different combinations of strategies. Every generation comprises interactions of all agents with all other agents in a repeated PD game ($T = 5$, $R = 3$, $P = 1$, $S = 0$; $p_s^* = 0.71$) with 200 repetitions. At the termination of each generation, the aggregated gains of each strategy are transformed into a new and proportional composition of agents, which then forms the population of agents for the next generation. Every population on the left side (panels A, C, and E) has a parallel population with an initially enlarged proportion of hostile and unpredictable agents presented on the right side (panels B, D, and F). Learning algorithms were equipped with memory spans of seven items; the MaRS passive and reactive similarity indices were activated after containing one and five items, respectively, and were restricted to a maximum of 10 items. Each panel consists of the averages over 100 evolutionary runs with identical initial conditions.

may comprise biased individuals with strong predispositions in regard to an opponent's similarity. Experiments testing human participants have already shown how individuals who undergo manipulations of perceived similarity respond in accordance with the SERS theory (10, 11). This mechanism highlights the negative impact of propaganda and prejudice, which alienate and distance opponents. It also calls for the development of simple intervention methods that promote cooperation among hostile populations by highlighting their similarity to each other.

Second, in the present simulation, we changed payoff structures of PD games without having to pay any cost. Clearly, this is not the case when trying to change the nature of economic, social or political conflicts. The logic of SERS underlying the function of MaRS explains the cooperative benefit of reduced similarity thresholds as well as the hostile potential of high similarity thresholds. This suggests that third parties who are willing to reduce conflict outbreaks, yet are not sufficiently powerful to completely change the nature of the encounter, should invest in reducing similarity thresholds. Based on the structure of the similarity threshold p_s^* (and the constraints of the PD game), we distinguish among two intervention types. The first, we term the robin hood approach because, like the legendary outlaw who robbed the rich and gave to the poor, this approach reduces the difference between the payoffs of the exploiter and the exploited. To this aim, one may reduce T, enlarge S, or, in accordance with Robin Hood's legend, transfer payoffs from T to S. The second, which we term the educational approach, enlarges the

difference between the payoffs obtained following mutual cooperation and the payoffs obtained following mutual defection. To this aim, one may enlarge R, reduce P, or transfer payoffs from P to R. Both approaches reduce the similarity threshold of the game, yet the educational approach is more efficient than the robin hood approach because it provides a more effective reduction for every extent of intervention (*SI Appendix, Appendix 2 and Fig. S3*, provide a formal proof).

Finally, we suggest the consideration of imitation and mimicry, such as those observed in chimpanzees (15) and young children (12–14), to provide the basic behavioral experiences needed for triggering the development of cooperative populations. Learning to mimic similar others, while also learning when to expect being mimicked by them, seems to provide a fundamental aspect of the evolution of cooperation.

ACKNOWLEDGMENTS. We thank the University of Bonn, the Max Planck Institute for Research on Collective Goods, and the University of Haifa research authority for providing infrastructure and support; Andreas Glöckner and the anonymous reviewers for many insightful comments; and Anton Gerstein for technical support. This work was supported by the German-Israeli Foundation for Scientific Research and Development Grant 1084-7.4/2010 (to I.F., S.J.G, and R.S.); Caesarea Edmond Benjamin de Rothschild Foundation Institute for Interdisciplinary Applications of Computer Science (A.F.); National Philanthropic Trust (S.L.); John Templeton Foundation Foundational Questions in Evolutionary Biology Program (S.L.); and NSF Grant IIS-0705822 (to D.I.R.).

- Knack S, Keefer P (1997) Does social capital have an economic payoff? A cross-country investigation. *Q J Econ* 112(4):1251–1288.
- Dugatkin A (1997) *Cooperation Among Animals* (Oxford Univ Press, Oxford).
- Wickler W (1968) *Mimicry in Plants and Animals* (World Univ Library, London).
- Payne RB (1997) *Host-Parasite Evolution: General Principles and Avian Models* (Oxford Univ Press, Oxford).
- Byrne DE (1971) *The Attraction Paradigm* (Academic, New York).
- Lakin JL, Chartrand TL (2003) Using nonconscious behavioral mimicry to create affiliation and rapport. *Psychol Sci* 14(4):334–339.
- Chartrand TL, Bargh JA (1999) The chameleon effect: The perception-behavior link and social interaction. *J Pers Soc Psychol* 76(6):893–910.
- Van Baaren RB, Holland RW, Steenaerts B, Van Knippenberg A (2003) Mimicry for money: Behavioral consequences of imitation. *J Exp Soc Psychol* 39:393–398.
- Cornelis I, Van Hiel A, De Cremer D (2011) Birds of a feather: Leader-follower similarity and procedural fairness effects on cooperation. *Eur J Work Organ Psychol* 20(3): 388–415.
- Fischer I (2009) Friend or foe: Subjective expected relative similarity as a determinant of cooperation. *J Exp Psychol Gen* 138(3):341–350.
- Fischer I (2012) Similarity or reciprocity? On the determinants of cooperation in similarity-sensitive games. *Psychol Inq* 23:48–54.
- Guillaume P (1926/1971) *Imitation in Children* (Univ Chicago Press, Chicago).
- Piaget J (1951/1967) *Play, Dreams and Imitation* (Routledge, London).
- Meltzoff AN, Moore MK (1977) Imitation of facial and manual gestures by human neonates. *Science* 198(4312):75–78.
- Whiten A (1998) Imitation of the sequential structure of actions by chimpanzees (Pan troglodytes). *J Comp Psychol* 112(3):270–281.
- Duersch P, Oechssler J, Schipper B (2012) Unbeatable imitation. *Games Econ Behav* 76:88–96.
- Apestequia J, Huck S, Oechssler J (2007) Imitation - theory and experimental evidence. *J Econ Theory* 136:217–235.
- Durrett R, Levin A (2005) Can stable social groups be maintained by homophilous imitation alone? *J Econ Behav Organ* 57:267–286.
- Axelrod R (1984) *The Evolution of Cooperation* (Basic, New York).
- Nowak M, Sigmund KA (1993) A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature* 364(6432):56–58.
- Miller GA (1956) The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychol Rev* 63(2):81–97.
- Skinner BF (1981) Selection by consequences. *Science* 213(4507):501–504.
- Harley CB (1981) Learning the evolutionarily stable strategy. *J Theor Biol* 89(4): 611–633.
- Osborne MJ, Rubinstein A (1998) Games with procedurally rational players. *Am Econ Rev* 88(4):834–847.
- Chmura T, Goerg SJ, Selten R (2012) Learning in experimental 2x2 games. *Games Econ Behav* 76(1):44–73.
- Selten R, Abbink K, Cox C (2005) Learning direction theory and the winner's curse. *Exp Econ* 8:5–20.
- Brown GW (1951) Iterative solution of games by fictitious play. *Activity Analysis of Production and Allocation*, ed Koopmans TC (Wiley, New York), pp 347–376.
- Berger U (2007) Brown's original fictitious play. *J Econ Theory* 135(1):572–578.
- Nowak A, Sigmund K (1992) Tit for tat in heterogeneous populations. *Nature* 355: 250–253.
- Press WH, Dyson FJ (2012) Iterated prisoner's dilemma contains strategies that dominate any evolutionary opponent. *Proc Natl Acad Sci USA* 109(26):10409–10413.