In a population with a local interaction structure, where individuals interact with their neighbors and learning is by way of imitating a successful neighbor, cooperation is shown to be a stable strategy that cannot easily be eliminated from the population.

Cooperation, Mimesis, and Local Interaction

ILAN ESHEL
Tel Aviv University
DOROTHEA K. HERREINER
University of Bonn
LARRY SAMUELSON
University of Wisconsin–Madison
EMILIA SANSONE
University of Naples
AVNER SHAKED
University of Bonn

How do modes of behavior spread in a society? What makes one convention more successful than another? In particular, how is it that cooperation and altruistic behavior are to be found so frequently in human societies, despite the obvious immediate disadvantage they cause?

Of the numerous explanations of this phenomenon that exist in the literature, two are often mentioned. The one, of a biological nature, due to Hamilton (1964), observes that within a group of genetically related individuals, behavior that makes individuals treat their kin kindly is more likely to succeed. Such behavior, if transmitted from generation to the next by genes, supports itself and propagates itself in the society. The cooperative behavior thus acquired by members of a family may have persisted when societies expanded beyond the immediate family group.

SOCIOLOGICAL METHODS & RESEARCH, Vol. 28 No. 3, February 2000 341-364

The other explanation, of a strategic nature, notes that cooperation is possible in situations with repeated interactions among the same, possibly unrelated, individuals. Behavior that punishes noncooperating individuals may be an equilibrium in the sense that all individuals will cooperate for fear of being punished. The punishing strategy that sustains cooperation requires long-term memory and ability to recognize a defecting individual.

We assume that individuals are boundedly rational, that they are not fully aware of the consequences of their actions or else are incapable of calculating the results of their actions. We assume that individuals use some rule of thumb, based on their current information and observations, to determine their next action. We offer an explanation for the existence of cooperation based on two simple assumptions. We assume that there is a local structure in the population; that is, individuals meet and interact only with their neighbors and not with the whole population. We also assume that individuals can observe the degree of success achieved by others and that they learn or update their behavior by imitating the behavior of their more successful neighbors.

Both assumptions seem plausible. Although the acquaintance network in a society is usually very intricate, it is only in very small societies that an individual may be familiar with all other individuals. Usually, individuals interact only with a small subset of the population: their neighbors, their friends, or their colleagues. Our second postulate, that individuals imitate behavior that has led to success, seems to be a cornerstone of human behavior and of our learning practices. Individuals, no doubt, attempt to use their analytic powers to understand a situation and find out in a rational way what is the best way to proceed. Unfortunately, many situations (indeed most) are too complicated to be fully analyzed with the individual's limited and bounded rationality. When the rational approach fails, people may resort to imitation or supplement their partial analysis by a dose of imitative behavior. Thus, even though some learning methods are more complex than mere imitation, imitation is the basis of many learning processes.

We have deliberately left the details of the local interaction system and the imitation process somewhat vague. We first present an intuitive argument why cooperation may survive under these assumptions and then fill in the missing details.

Our two assumptions (local interaction and imitation) are shown to lead to the survival of cooperative behavior. We offer here intuitive arguments that will be made precise in the following sections. To clarify our point, assume that individuals are located on a line, each interacting with and imitating those located within a certain distance from him or her. Imagine that an individual can be either cooperative and helpful to his neighbors or egoistic and noncooperative. Consider first an isolated noncooperating individual surrounded by cooperative players. He does well, for his neighbors support him while he makes no effort to help them or perhaps even exploits them. He will therefore be imitated by his neighbors, this will create a group of noncooperating individuals that is no longer doing well, since no one helps his or her neighbors. So an isolated island of noncooperation will tend to spread but not too much. Once too big, it no longer is an object for imitation. With cooperation, the dynamics seem to go the other way. An isolated cooperator will not do well if surrounded by noncooperators. He helps his neighbors who do not help him. His neighbors, on the other hand, benefit from his being there. His mere existence makes noncooperation more attractive in his neighborhood, and he himself will eventually imitate his neighbors. In contrast, a large group of cooperators is strong in the sense that its members support each other and are therefore (locally) successful. If the cooperators on the group's boundary (where they interact with and observe noncooperators) hold fast, then the cooperative group will not be eroded and may even spread.

This intuitive argument seems to suggest that if cooperation survives, it will be present in large numbers interspersed with small islands of noncooperation. The intuitive argument leaves open the question of whether cooperation can survive at all. To find out whether cooperation survives, we need to analyze the problem in detail, for the survival or extinction of cooperation crucially depends on what happens at the boundaries between groups of cooperators and noncooperators. This question becomes even more elaborate once the learning process is allowed to be stochastic (the individual chosen to be imitated is chosen at random) or when noise is introduced to the system in the form of mutations. The models we introduce in the following sections address this problem and establish conditions under which cooperation will survive.

MODEL 1: ALTRUISTS AND EGOISTS[1]

THE MODEL

Assume a finite population of individuals located on a circle. Let there be two types of individuals: altruists and egoists. An altruist provides a public good to his neighbors at a (net) cost to himself. An egoist enjoys, at no cost, the public goods provided by his altruist neighbors while giving nothing in return. To pin matters down, assume that an individual's neighborhood consists of his *two* immediate neighbors on the circle and that an altruist increases the utility of each of his two immediate neighbors by one unit at a negative net benefit of $-C$ units to himself,[2] with $0 < C < 1/2$. We further assume that the total benefit of an altruist is the number of altruists among his neighbors minus C and that the benefit of an egoist is the number of altruists in his immediate neighborhood.

The learning-imitation process is as follows. Let each individual be informed of the average payoff of each type in his immediate neighborhood (consisting of himself and his two neighbors). On the basis of this information, the individual will switch to the strategy (type) that earns the higher average payoff. If an individual is surrounded by identical types, he does not change his behavior. This learning process describes a situation in which an individual does not observe other individuals in detail but receives general information about his environment and how on average each strategy fares there; he also believes that this information is relevant to his future success.

Clearly, our individuals are not aware that they are, in fact, playing a version of the Prisoners' Dilemma. If they understood it, then in a one-shot game they would all become egoists, since it is always better to be an egoist whoever your neighbors are. The simple situation described above should be understood to represent a complicated one in which it is difficult to figure out what the benefits of each action are.

We now have a dynamic process. Beginning at an initial state, a configuration of altruists and egoists on the circle, each obtains a payoff, learns about the payoffs in his neighborhood, and chooses the type he will be in the following period, thus bringing the process to a new state and so on.

THE LIMIT OF THE DYNAMIC PROCESS

The process defined above, being a deterministic Markov process on a finite space, converges to an absorbing set: a set of states in which it cycles and from which it never exits.

It can easily be shown that beginning with any initial configuration of altruists and egoists on the circle, the process converges to one of the following states or cycles:[3]

1. A state in which all individuals are altruists
2. A state in which all individuals are egoists
3. A steady pair of egoists: A state in which all but two adjacent individuals are altruists (an altruist is represented by a and an egoist by E)

   ... aaaaaaaaEEaaaaaaa ...

4. A blinker: A cycle of two states in which egoism spreads from an isolated egoist to a string of three egoists and shrinks again because as a larger group egoists are not doing well:

   ... aaaaaaaaEaaaaaaa ...

   ... aaaaaaaEEEaaaaaa ...

   (The first line describes the situation at a given time and the next line describes the new configuration at the following period.)
5. A cycle of two states in which isolated strings of egoists (blinkers and steady pairs of egoists) exist among altruists; for example,

   ... aaaaaaaaEaaaaEEaaaaEEEaaaaa ...

   ... aaaaaaaEEEaaaEEaaaaaEaaaaaa ...

It is easy to verify that in an absorbing state of type 3, no individual wishes to switch to another strategy, and that in a blinker (an absorbing state of type 4), only the immediate neighbors of the egoist at the center wish to change their type and do so in a cyclical manner.

Note that in absorbing sets of type 5, the islands of egoists need be well apart from each other. If they happen to be too close, as in the following example, the islands merge to create a large group of egoists that is no longer viable and will therefore shrink, ending with a single blinker and fewer egoists than we had initially:

```
. . . aaaEaaEEEaaEaaaaaa . . .

. . . aaEEEEaEaEEEEaaaaa . . .

. . . aaaEEEEEEEEEaaaaaa . . .

. . . aaaaEEEEEEEaaaaaaa . . .

. . . aaaaaEEEEEaaaaaaaa . . .

. . . aaaaaaEEEaaaaaaaaa . . .

. . . aaaaaaaEaaaaaaaaaa . . .

. . . aaaaaaEEEaaaaaaaaa . . .

. . . aaaaaaaEaaaaaaaaaa . . .
```

The minimal number of altruists between strings of egoists can be calculated, and it can be shown that in any absorbing set where altruists are present, there will be at least 60 percent altruists. Thus, an absorbing set has either *no* altruists or *a majority* of altruists. This seems to confirm our intuition that altruism, if it survives, will be present in large numbers. However, we need to check whether and how often altruism survives at all.

The other part of the intuitive argument, that a clump of altruists supports itself and may expand, can be seen by the following example:

```
. . . EEEEEEaaaaaEEEEEEEEE . . .

. . . EEEEEaaaaaaaEEEEEEEE . . .

. . . EEEEaaaaaaaaaEEEEEEE . . .

. . . EEEaaaaaaaaaaaEEEEEE . . .
```

A sufficiently large string of altruists (here a string of five) expands in a sea of egoists. Moreover, it can be shown that a string of five or more altruists is never destroyed; it need not expand as it does in the above example, but it will never be eradicated.

Where we end up depends on where we start: The absorbing set to which the process converges depends on the initial configuration of altruists and egoists on the circle. There are 2^N such configurations, where N is the length of the circle. We can now ask how many of the initial configurations lead to an absorbing set in which there is *no* altruism. Combinatorial arguments demonstrate that for large N, nearly all of the 2^N possible configurations will contain a string of (at least) five altruists. Because such a string never vanishes, it follows that most of the initial configurations lead to an absorbing state in which altruists exist and are therefore a majority of at least 60 percent. Thus, for a sufficiently large population on a circle and beginning with an arbitrary initial configuration, it is with probability (approaching) 1 that the absorbing set we reach will have a majority of altruists. This confirms our intuition that in this model, altruism is very likely to survive in large numbers in the population.

So far, we have assumed that the imitation process is deterministic and fully synchronized: All individuals update their strategies each period. These assumptions may be relaxed without changing the result. Imagine that an individual observes the payoffs of his neighbors, as before, except that now, instead of switching *with certainty* to the strategy with the higher average payoff, he does so with a positive probability that is not necessarily 1. Note that an individual always revises his strategy to a "better" one; he may not switch to a strategy that earns a lower average payoff than the one he currently uses. Choosing not to switch to a better strategy when there is one available effectively means forgoing the opportunity to learn. Thus, under this learning scheme, learning is no longer synchronized.

It is interesting to see that under this stochastic learning process, all the above results are valid. The only difference is that a blinker (type 4) is no longer part of an absorbing set. A blinker requires that both neighbors of an isolated egoist simultaneously update their strategy; this is no longer the case in stochastic learning. However, steady strings of egoists (type 3) are still stable. As for deterministic learning, nearly all initial configurations lead to absorbing sets consisting of at least 60 percent of altruists.

MUTATIONS

So far, we have assumed that the individuals rigidly follow the learning rules and switch only to a strategy that earns a higher average payoff than the one they are currently using. It is unreasonable to assume that deviations from this procedure do not occur. Mistakes happen and individuals may sometimes act erratically. To this purpose, mutations are introduced. After learning, an individual may have a small probability of becoming a mutant and determine his future type arbitrarily. Mutants introduce noise to the system. With mutants, it is no longer true that the state in which all are altruists is a stable state: Some individuals may mutate to become egoists, and once egoists they may have no reason to switch back to altruism.

On the face of it, it seems that mutations make it more difficult for altruism to survive. Mutations seem to increase the number of egoists: If a mutant egoist appears in a population of altruists, it will do well and survive, whereas if a mutant altruist appears in a population of egoists, it will be eradicated. It may therefore come as a surprise that altruism survives mutations as long as these occur with a sufficiently small probability or as long as the circle (N) is sufficiently large. The mathematical details of this result are rather intricate and can be found in Eshel, Samuelson, and Shaked (1998). Here, we present an intuitive argument that drives the result.[4]

If mutations are rare, then the dynamics will, most of the time, follow the learning process. After a sufficiently long time in which no mutations occur, the process will settle in an absorbing set of the learning procedure. When mutations appear, they will disturb the dynamics and may kick the population to a state outside this absorbing set (or its basin of attraction). Again, after a sufficiently long time, the population will end up in (possibly) another absorbing set. As we have seen above, there are two types of absorbing sets: those that have at least 60 percent altruists and a single state in which there are no altruists. Because we are only interested in the existence of altruists (and not in their precise distribution on the circle), we need only find out how easy it is for mutations to shake the population away from one type of absorbing set to the other; that is, is it easier for mutations to create persistent groups of altruists when there are none than to eliminate all altruists when they are the majority?

It is relatively easy for mutations to create a string of five altruists in the state where all are egoists; it requires only five simultaneous mutations. It is not easy for mutations to eradicate *all* strings of altruists when at least 60 percent of the population is altruistic. It will not do to destroy the strings one by one as our example in the previous section shows. Destroying a string of altruists between two strings of egoists creates a large string of egoists. This large string is not viable and will shrink to a small string of egoists, ending up with more altruists than there were initially. To completely destroy the altruists, all strings of altruists need to be simultaneously eradicated. Depending on how many strings of egoists there are, the strings of altruists are either numerous or long. It can be shown that to destroy all of them, one needs to introduce simultaneously a sufficient number of mutations in each string of altruists. To do this, it can be shown that at least $N/10$ mutations are required. When N is sufficiently large, it is much larger than five, the number of mutations needed to move in the other direction, and so the population will be most of the time absorbing sets with a majority of altruists. Thus, even noise in the form of mutations cannot eliminate altruism.

ROBUSTNESS AND FRAGILITY OF THE MODEL

Our model has the following basic components: a population located on a circle where each individual can be one of two types, altruist or egoist, and each interacts and learns from his two immediate neighbors. The intuition presented in the introduction suggests that the result (that altruism will be significantly present in the population) will hold even when the model's basic features will be slightly changed. In this section, we present some generalizations of the model.

The method of analysis we developed can be applied to populations in which the interaction between individuals takes the form of a general normal form game and a type is characterized by the strategy he plays. We can analyze all 2×2 games; that is, we can follow the evolution of any two strategies in the population and describe some properties of the limit distribution, as we have for altruists and egoists. In particular, altruists and egoists can be viewed as a special case of a type who contributes K to each of his neighbors at a net benefit of $-C$ to himself. An egoist has $K = C = 0$, whereas an altruist has $K = 1$, $C < 1/2$. Our model can be extended to any two or indeed more such types. For example, one may consider a hooligan, an individual who reduces the utility of his neighbors at a possible benefit for himself, $K = -1$, $C < 0$. Compared to a hooligan, an egoists is a (relative) altruist, for he at least causes no damage. It can be shown that in a population of egoists and hooligans, the egoists will eventually be a majority, with small pockets of hooligans between them.

In our model, we have assumed that both neighborhoods, the interaction neighborhood and the learning neighborhood, are of radius 1. Our method enables us to analyze the dynamics when the neighborhoods are of a larger radius. Here, the cost of being an altruist plays an important role, and the percentage of altruists in the population is a function of that cost.

We have not been able to obtain analytic results for a population located on a two-dimensional grid. A number of works have attempted to find regularity in the evolution of two-dimensional populations using computer simulations (Nowak and May 1992, 1993; Nowak, May, and Sigmund 1995).

Our analysis becomes difficult when the neighborhoods are large (with radius >1) and the learning process stochastic. In this case, an individual may by learning import a strategy from far away, and new strings of egoists may emerge. Also, the existing strings expand and contract at random. All this makes it very complicated to analyze the dynamics of such a population. Instead, we simply want to know which strategies can withstand an invasion of mutants. This we do in Model 2.

MODEL 2: CONSERVATIVE STOCHASTIC LEARNING AND UNBEATABLE STRATEGIES[5]

This model features an infinite population located on a line. Each individual, when called to revise his strategy, chooses a strategy from his neighborhood according to some probability. The probabilities are such that he is more likely to imitate a successful individual in his neighborhood. In contrast to the previous model, he may choose a strategy that is "less successful" than his current one.

Let individuals be located at the integer points of an infinite line (0, ±1, ±2, ±3, . . .). The interaction between players takes the form of a Prisoners' Dilemma,[6] with $c > a > d > b > 0$ (see Table 1).

An individual has a strategy (mixed or pure) that he plays against all neighbors in his interaction neighborhood, which consists of k indi-

TABLE 1

	C	D
C	a,a	b,c
D	c,b	d,d

viduals to his right and k to his left. After obtaining their payoffs, a single individual is chosen at random to learn and update his strategy. He does so by picking up a "guru" in his learning neighborhood, a neighborhood of radius n around him, whose strategy he will imitate. Choosing a guru is a probabilistic process. A guru is chosen with a probability that is his relative success in that neighborhood; that is, if the total payoff of the individuals in the learning neighborhood (including the learning individual himself) is M and a particular individual has the payoff m, then the probability that he will be chosen to be the guru is m/M.[7] For technical reasons, time in this model is taken to be continuous, so that at each point in time individuals earn a payoff, and very few individuals are chosen, according to a Poisson process, to learn and update their strategy. The population is taken to be infinite so that even large neighborhoods (large n,k) will be local relative to the population.

We introduce an additional assumption about learning that distinguishes between cultural and biological evolution. We assume that individuals are conservative and are not keen to adopt strategies that are not popular in their neighborhood. Thus, when called to learn, an individual will not change his strategy if his two immediate neighbors play the same strategy as he does. Only if at least one of his immediate neighbors plays a strategy different than his—that is, he is on a border between regions of two strategies—will he look at his learning neighborhood (of radius n) and choose an object for imitation.

This conservative learning process is a stringent version of a learning process that requires an incentive to learn by way of gradual search. When called upon to learn, a player will consider his two immediate neighbors; if neither is different than him, he will with high probability abandon his search and give up learning. With small probability, he will continue to search and look at his four immediate neighbors; if none of those plays a strategy different than his, he will with high probability abandon learning. With small probability, he will continue searching for an incentive to learn among his six immediate neighbors, and so on. The process we have chosen is a simplified version of this process in which a player stops his search if he does not receive an immediate incentive to learn.

We do not claim that cultural learning is always conservative, only that it sometimes is, and that it is not to be found in biological evolution. In a similar biological process, an organism occasionally dies and is replaced by a clone of one of the organisms in his learning neighborhood. Here, the learning neighborhood represents how far an organism can shoot his seed. A more successful organism is more likely to plant his seed in the vacant location. This will be done irrespective of whether the dead organism was similar to his neighbors. Hesitation to adopt a strategy different from the one prevailing in one's neighborhood is a property of cultural and not biological evolution. In the Relaxing Conservatism section, we show that the two processes representing biological and cultural evolution may lead to different results.

The path of such an elaborate stochastic process is difficult to analyze; instead, we concentrate on a different aspect of the evolution. We ask whether there exists a strategy that is immune to an invasion of mutants in the following sense: Is there a strategy that when all individuals on the line play it and a finite number of identical mutants appear, the stochastic learning process will eliminate the mutants with probability 1? We call such a strategy, if it exists, an unbeatable strategy.[8] The readers familiar with the biological and game theoretic literature will notice the similarity of this concept to the concept of an evolutionarily stable strategy (ESS).[9]

Our conservative learning ensures that the mutants remain concentrated in the population. Because only individuals on the boundary between the mutants and the indigenous population may learn, and by doing so may shift the border between the two strategies, no new strings of mutants will be created by the process.

UNBEATABLE STRATEGY IN A POPULATION PLAYING THE PRISONERS' DILEMMA

In a population playing the Prisoners' Dilemma, the unbeatable strategy depends on the ratio $\theta = n/k$ between the radii of the learning and interaction neighborhoods. When θ is large, it is the cooperative strategy of the Prisoners' Dilemma that is unbeatable. This happens when an individual interacts only with people in his village but learns from a large neighborhood by reading a national newspaper; in such situations, new ideas may travel far when adopted by faraway people. When θ is small, in the unlikely case when individuals interact with a large group but adopt ideas only from their close neighborhood, then it is the noncooperative defect strategy that is unbeatable.

This result is due to the conservative learning process. Imagine a situation in which all individuals play the cooperative strategy C and a finite number of mutants playing another strategy invade the population. The mutants' strategy, being different from the cooperative strategy, is necessarily less cooperative; that is, it gives a mutant who confronts a mutant a payoff less than a (the payoff of a cooperator playing against a cooperator). The finite number of mutants on the infinite line guarantees that the mutants have only a local effect; most individuals on the line do not confront the mutants at all. The conservative learning process ensures that only individuals on the boundaries between cooperators and mutants may learn and possibly change their strategy, thereby shifting the boundary one step to the right or to the left. Mutants remain therefore concentrated on the line; no new strings of mutants appear, and only the existing strings expand or shrink. It can be shown that it is sufficient to study situations in which the mutant strings have expanded to form large patches and test whether in these situations mutants continue to expand. When θ is large (i.e., n is very large compared to k), an individual on the boundary who is permitted to learn will look deep into the region where cooperating individuals interact with their own kind (since the radius of their interacting neighborhood k is relatively small) and are therefore doing very well. On the other hand, an individual may observe many mutants playing against mutants, whose payoff is therefore less. Thus, he is more likely to choose a cooperator as his guru and switch to cooperation. The group of mutants is therefore likely to shrink, and it can be shown that it will be eliminated with probability 1. This argument is reminiscent of the intuitive argument presented in the introduction: Defection cannot spread too much, for then the individuals on its boundary will observe prosperous cooperators and imitate them.

For small θ (where n is considerably smaller than k), consider a population playing the defect strategy, with only a finite number of mutants playing another strategy. An individual about to learn considers imitating players who sit close to him (imagine n to be small). Those individuals interact with a large neighborhood (imagine k to be large). The differences between the various interaction neighborhoods of the few individuals in the learning neighborhood are therefore insignificant. All the observed individuals face, practically, the same interaction environment. Against a *given* neighborhood, defection always earns the highest payoff. Thus, a learning individual is likely to become a defector and the mutant group of cooperators will shrink. This ensures that defection is the unbeatable strategy when θ is small.

This result indicates that when learning is conservative (in cultural evolution), a high mobility of ideas leads to cooperation. This result will be shown to depend crucially on the conservatism of the learning process.

UNBEATABLE STRATEGY IN A POPULATION PLAYING AN ARBITRARY GAME

In the case of the Prisoners' Dilemma, it was the nature of the two strategies in the game, together with the ratio $\theta = n/k$, that enabled us to determine the unbeatable strategy. In a general game, this is more complicated, since the strategies do not necessarily have the straightforward properties of those in the Prisoners' Dilemma. First, assume that the game has a unique strategy x that maximizes the payoff when playing against itself (the cooperative strategy in the Prisoners' Dilemma has this property). This strategy is successful against itself; however, there is no guarantee that it does well against other strategies. For a sufficiently large n (holding k fixed), this strategy x will be the unbeatable strategy. The argument is similar to the one in the previous section. Let all the population play strategy x and let a finite number of mutants play another strategy. A player on the boundary between the mutants and the indigenous population will see mostly individuals playing against a strategy identical to their own, mutants playing against mutants and x players playing against x players. By the choice of x, the x players do better than the mutants. He sees a few players who confront a strategy different than their own, but this does not change the fact that it is the x players who are successful in his neighborhood, and that he is therefore likely to become an x player himself. Thus, the x strategy will take over and eliminate the mutants.

In the case when k is not small, or when the ratio between n and k is not small, the individuals close to the boundary who play against strategies different than their own are no longer an insignificant group and may influence the learning individual's decision. In that case, a more subtle argument is required: The candidate strategy for unbeatability should do well against itself but at the same time it should do well against other strategies. The balance between these two requirements depends on how many of the x players that the learner observes interact with strategy x and how many with the mutants. If most of the interactions of those x players are with x, then in order for x to be unbeatable it should do well against itself. If most of the interactions of the observed x players are with mutants, then x should do well against any mutant and so better be aggressive if it is to eliminate any mutant. Thus, if most of the interactions are "within the family," then x should be friendly to itself, whereas if its interactions are mostly with outsiders, then what counts is that it should do well against them. It therefore seems as if an unbeatable strategy measures whether it is related to its opponents, and if it is, it plays cooperatively. This idea is captured by Hamilton's (1964) inclusive fitness. Hamilton suggested that when individuals are related to a degree r, then the relevant payoff of an interaction should be given by the individual's payoff plus r times the payoff of his related opponent. Thus, when the opponent is my twin brother (i.e., $r = 1$), I give his welfare the same weight as to mine. Given the parameter r, a new game is derived from the interaction game in which the payoff to a player of an interaction consists of his original payoff plus r times his opponent's. This degree of "kinship" r, which in our case measures how often an x player interacts with similar x players (as opposed to mutants), is a simple increasing function of $\theta = n/k$.

An unbeatable strategy can be shown to be a Nash equilibrium in this new game. In fact, an unbeatable strategy is more than a Nash equilibrium; it is an ESS of this inclusive fitness game.

The intuition emerging from this analysis is rather surprising: A strategy that is "strong" in the population with the local interaction structure is found to be one that is "friendly" to similar opponents. It is as if kinship behavior arises from frequent local imitations. There is no need for the individuals to be related; the frequent local interactions guarantee that they will behave as if they were related.

RELAXING CONSERVATISM

The analysis of the theoretical models presented in the previous section relies on the assumption of conservative learning. A prompt incentive is needed to make an individual revise his strategy; he needs to see a strategy different than his own between his two immediate neighbors to consider changing his strategy. In this section, we test whether the results we obtained are robust to changes in this assumption.

The evolutionary process may be interpreted in two ways, as cultural learning or as a biological propagation process. The way to relax the conservative learning assumption depends on which interpretation we adopt. If it is cultural evolution that we are interested in, then making the individual less conservative should mean that he may learn even if his two immediate neighbors play strategies identical to his, but he should be reluctant to do so. We capture this idea in the following way. An individual will learn with certainty when one of his immediate neighbors plays a strategy different than his. However, he will be permitted to learn with a certain probability when his two immediate neighbors play a strategy identical to his but one of the next two neighbors plays a different strategy. This means that an individual will learn with certainty when he is located on the border between two strategies and will only sometimes learn when he is not directly on a border. Note that this probability governs only his incentive to learn; once willing to learn, he observes and considers a neighborhood of radius n.

We expect that for small values of this probability, a strategy that was unbeatable with conservative learning will remain unbeatable. For, as before, learning will be mostly conservative; only with small probability will an individual introduce a new strategy to his immediate neighborhood.

TABLE 2

	C	D
C	S,S	$S+1,1$
D	$1,S+1$	$2,2$

If, however, evolution should describe biological propagation, then no such conservative restriction is appropriate. An individual dies and is replaced by an offspring of a neighboring individual. This replacement is independent of whether or not the deceased individual was identical to his two immediate neighbors. Thus, in the biological case, relaxing the conservative assumption would mean that an individual can "learn" whenever someone in his learning-propagation neighborhood (a neighborhood of radius n) plays a strategy different than his own.

In this case, we do not necessarily expect the results of the theoretical model to hold. An invading mutant may take over a population playing a strategy that is unbeatable in the previous model. The fact that a mutant can spread far from his current location due to the new learning rule, and that this may happen often (in contrast to cultural learning, where this rarely happens), makes the existing population vulnerable to invasions by mutants.

Both ways of easing the learning rules (the cultural and the biological) lead to models that are analytically intractable. We therefore turn to computer simulations to test our intuitions.

COMPUTER SIMULATIONS

For the computer simulations, we consider the Prisoners' Dilemma shown in Table 2 (with $S > 2$). The first strategy C (cooperate) is an unbeatable strategy when the radius of the learning neighborhood is sufficiently large relative to the radius of the interaction neighborhood: $n >> k$.

Throughout this section, we keep the interaction neighborhood small, $k = 1$, and the learning neighborhood relatively large, $n = 3$. The parameter S, which measures the cooperator's contribution to his opponent, took the values 3, 5, and 7; later, it was increased to $S = 18$. It is straightforward to check that for the above game (and for $k = 1$, $n = 3$), the strategy C is unbeatable for $S > 2.4$. (Increasing S makes it more advantageous for a cooperator to be surrounded by cooperators; similarly, increasing n makes C more robust against D mutants.)

We wished to test whether C remains unbeatable when the learning rule changes. The learning rule was changed in two ways, describing cultural and biological evolution. To move smoothly between the two, we introduced a variable p, $3 \geq p \geq 1$. For $2 \geq p \geq 1$, the individuals on the boundary learn with certainty when given the opportunity whereas their immediate two neighbors learn with probability $p - 1$ when asked to learn. For $3 \geq p \geq 2$, the four individuals close to a boundary learn with certainty whereas the next two individuals learn with probability $p - 2$. Because $n = 3$, there was no need to consider higher values of p. Thus, p measures the divergence from conservatism. For $p = 1$, the learning is conservative. For values less than 2, the learning is cultural, and players not located on a boundary are reluctant to learn. For higher values of p, the model describes biological evolution. In the simulations, p was given values between 1 and 3 in increments of 0.1.

Unbeatability of cooperation when learning is conservative implies the two following properties, which were tested in simulations for learning processes for which the conservative assumption is relaxed:

1. When all the population plays defect except for a sufficiently large interval of players who play cooperate, the cooperators should eliminate the defectors.
2. When all the population plays cooperate and mutants playing defect enter, the mutants should be eliminated with certainty.

In addition, we checked whether increased values of S can help cooperation beat mutant defectors even in the biological evolution model, where there is no conservatism ($p = 3$).

For the simulations, we took a population of 1,000 individuals located on a line. Of those individuals, most played the same indigenous strategy, whereas mutants were introduced in some center positions of the interval. The strategy cooperate was sometimes the indigenous strategy and sometimes the mutant strategy. In both cases, we were interested in whether cooperate took over the entire population.

The simulations were run as follows. For each pair of values p,S and beginning with some configuration of cooperation and defection on the line, we first calculated the average payoff of each individual. We then randomly chose an individual to learn and tested whether he was permitted to learn; if so, he will have used the stochastic learning rule and imitated one of his neighbors.[10] This led to a new configuration for which we calculated the payoffs and so on. We continued the process until one of the strategies took over the entire population. A time limit of 10^7 periods that was incorporated in the program was never met. For each pair of p,S we iterated this experiment 1,000 times and noted in how many of those runs the strategy cooperate won the population over.

In the first set of simulations, all but the central 10 individuals played defect whereas the 10 mutants played cooperate. The block of 10 proved to be sufficiently large so as to allow the mutants (cooperate) to win when learning was close to conservative. In the second set of simulations, the roles of cooperate and defect were swapped; then, all but a few mutants were cooperators. Here as before, we asked whether cooperation won; to obtain convincing results, we made it as difficult as possible for cooperation to win. Defectors are better off when they are isolated; hence, we introduced four isolated defectors in the center of the population (in positions 1, 7, 13, and 20 of the central 20 positions).

SIMULATION RESULTS

In the first simulation, a block of mutant cooperators was introduced in a population of defectors. Figure 1 clearly shows that for small values of p, when learning was close to conservative, cooperation won the population over. For higher values of S, when cooperation had a greater advantage, cooperation won close to 100 percent of the runs even when departure from conservatism was significant. For lower values of S, the curve was not as flat; that is, the results are sensitive to deviation from conservatism. Note that below $S = 2.4$, cooperation is no longer unbeatable; for $S = 3$ when cooperation is close to being beatable, the block of 10 cooperator mutants may not have been sufficiently large to enable them to win in all cases when learning was conservative, but they still won more than 75 percent of the runs. When p was close to 3, and the model was more adapted to biological evolution, defection comfortably won the population over.

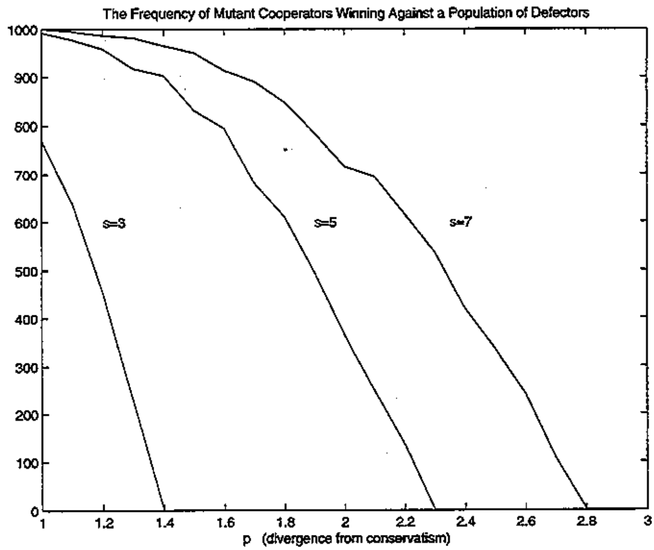

Figure 1: Test 1

In the second simulation, four defector mutants were placed in a population of cooperators. Figure 2 shows that cooperation eliminated the mutants unless the deviation from conservatism became significant. The results are not as clear at the "biological" end of the figure, where p is close to 3. For $S = 7$, cooperation continued to win in most runs.

The results seem to suggest that as S increases, cooperation can eliminate mutant defectors. This is verified by the simulations depicted in Figure 3. Here, we increased S to $S = 14$ and found that for high values of S, cooperation eliminated defection in almost all the runs even when learning was not conservative.

SUMMARY AND CONCLUSIONS

We studied the effects of a local interaction and learning structure on the evolution of a population in which individuals learn by imitating a more successful neighbor.

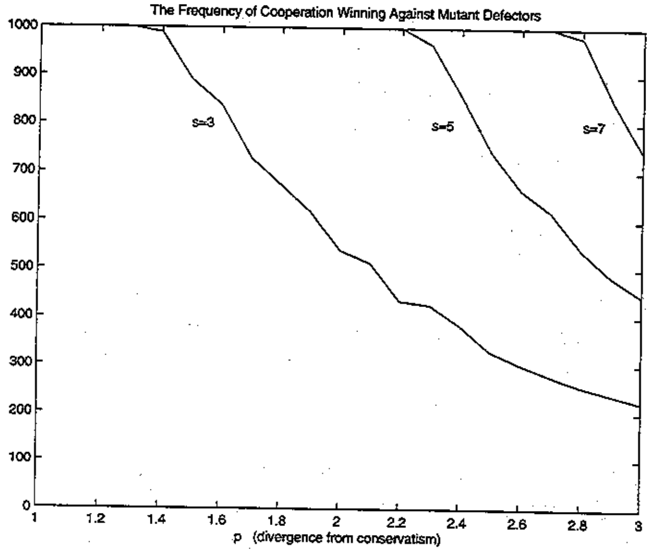

Figure 2: Test 2

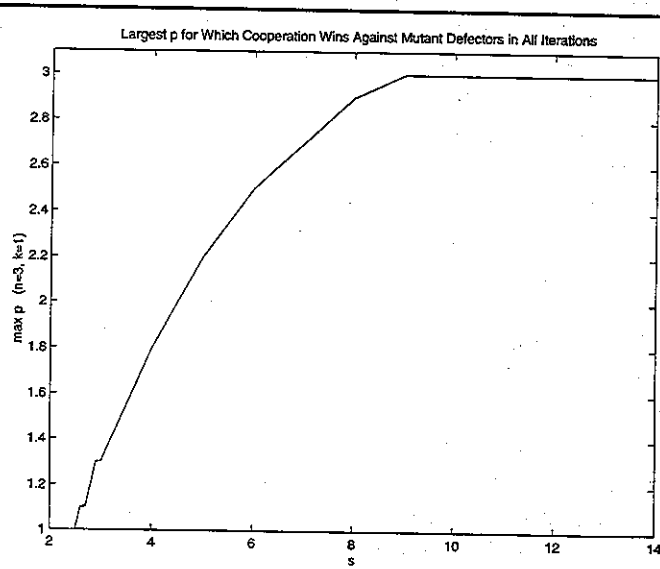

Figure 3: Test 3

We presented two models. In the first model, a finite population was located on a circle, and each individual learned by switching to the strategy earning the highest average payoff in his neighborhood. We found that when the circle is sufficiently large, with high probability the population will evolve to have a majority of cooperators. We found that even the introduction of mutations cannot destroy the robustness of cooperation in large populations. However, the combination of large learning neighborhoods with stochastic learning proved to be too much for sustaining cooperation. It enabled a single defector to spread and take over the (initially) wholly cooperative population. Although this requires a chain of events that occurs with low probability, it suffices to make the model analytically intractable.

We therefore introduced the second model in which we tested, as in the biologically inspired ESS, the stability of certain configurations. Here, learning was done by probabilistically imitating one of the neighbors. The probability of imitating someone is proportional to his relative success in the learning neighborhood. In this model, we looked for an unbeatable strategy, a strategy that if all play it and a finite number of mutants enter, the learning process will eliminate them with probability 1. The unbeatable strategy, when it exists, can be shown to be an ESS of a related interaction game in which each individual cares about the payoff of his opponent to a degree determined by the radii of the learning and interaction neighborhoods.

Thus, the message of this model is that cooperation could have developed as a result of the local nature of interaction and learning. To obtain this result, we assumed that learning is conservative: Individuals do not learn when their immediate environment does not provide them with incentives to learn. We argue that conservative learning may be a feature of some cultural evolutionary processes but not of biological evolution.

We investigated by computer simulations whether our results are robust when the conservative learning assumption is relaxed. We examined a situation in which cooperation is unbeatable (with the conservative assumption) and tested whether it remains unbeatable when conservatism is relaxed. We found that as we moved away from cultural toward biological evolution, cooperation was no longer unbeatable and defection could win the population over. However, cooperation can survive even in a biological process when the altruistic act costs little relative to its contribution to the beneficiary. Thus, local interaction supports cooperation against invading new forms of behavior in conservative cultural environments and even in biological setups when the costs of helping the other are sufficiently small.

NOTES

1. This model is described in detail in Eshel, Samuelson, and Shaked (1998). A similar model can be found in Bergstrom and Stark (1993).

2. We assume that his costs of providing the public good are $1 + C$, so that his net benefit is negative: $1 - (1 + C) = -C$.

3. It is our assumption $C < 1/2$ that enables the long-term existence of altruists. If $C > 1/2$, it is so disadvantageous to be an altruist that altruists can no longer coexist with egoists. In that case, the only absorbing state that includes egoists is the one in which there are no altruists.

4. Rare mutations were first introduced to evolutionary game theory by Young (1993) and Kandori, Mailath, and Rob (1993).

5. The details of this model are taken from Eshel, Sansone, and Shaked (1999).

6. We will see what happens when the game is not the Prisoners' Dilemma.

7. The results of this model remain unchanged when the probabilities of choosing the next strategy depend on the average payoff of each strategy (as in the previous model of altruists and egoists) and not on the sum of payoffs.

8. An unbeatable strategy has the additional feature that if it appears as a mutant in a population playing another strategy, then it has a positive probability of defeating the other strategy and winning the population over.

9. The difference between the concepts is that unbeatability is defined for a population with a local interaction structure whereas evolutionary stable strategy (ESS) is defined for totally mixed, panmictic populations. Moreover, ESS does not specify the process by which the mutants may be eliminated, whereas we have a well-defined process.

10. Both payoff and learning of the individuals situated at the edge of the population were modified to account for their fewer neighbors. This cannot be avoided in a finite simulation but has little influence on the results.

REFERENCES

Bergstrom, Theodore C., and Oded Stark. 1993. "How Altruism Can Prevail in an Evolutionary Environment." *American Economic Review* 83:149-55.

Eshel, Ilan, Larry Samuelson, and Avner Shaked. 1998. "Altruists, Egoists and Hooligans in a Local Interaction Model." *American Economic Review* 88:157-79.

Eshel, Ilan, Emilia Sansone, and Avner Shaked. 1999. "The Emergence of Kinship Behavior in Structured Populations of Unrelated Individuals." *International Journal of Game Theory* 28:447-63.

Hamilton, William D. 1964. "The Genetic Evolution of Social Behavior." *Journal of Theoretical Biology* 7:1-52.

Kandori, Michihiro, George J. Mailath, and Rafael Rob. 1993. "Learning, Mutation, and Long Run Equilibria in Games." *Econometrica* 61:29-56.

Nowak, Martin A., and Robert M. May. 1992. "Evolutionary Games and Spatial Chaos." *Nature* 359:826-29.

———. 1993. "The Spatial Dilemmas of Evolution." *International Journal of Bifurcation and Chaos* 3:35-78.

Nowak, Martin A., Robert M. May, and Karl Sigmund. 1995. "The Arithmetics of Mutual Help." *Scientific American* 272 (6): 76-83.

Young, Peyton. 1993. "The Evolution of Conventions." *Econometrica* 61:57-84.

Ilan Eshel is a professor of biomathematics at Tel Aviv University. His research interests are theoretical population biology, long-term dynamics of complex genetic systems, population, game theory, and complex dynamics in structured populations.

Dorothea K. Herreiner is an assistant professor in the Department of Economics at the University of Bonn. Her research interests include decentralized market models, local interaction, learning, matching, and electronic money.

Larry Samuelson is a professor of economics at the University of Wisconsin–Madison. His research interests center on game theory, particularly the use of evolutionary models to incorporate ideas of bounded rationality and learning.

Emilia Sansone is a researcher of mathematical physics at the University of Naples. Her research interests include mathematical physics, long-term dynamics of complex genetic systems, population game theory, and complex dynamics in structured populations.

Avner Shaked is a professor of economic theory at the University of Bonn. His research interests are game theory and evolutionary models of economics, in particular local interaction and learning.