

Norms of Cooperation*

Cristina Bicchieri

REDUCTIONIST STRATEGIES

A longstanding tradition in the social sciences contrasts instrumental rationality and social norms as alternative ways of explaining action. Rational choice theory defines action as the outcome of a practical inference that takes preferences and beliefs as premises. An explanation in terms of norms depicts a socialized actor whose behavior is not outcome oriented, since when acting in accordance with a norm one does not engage in a rational calculation nor does one pay very much attention to the consequences. Attempts at bridging the gap have either tried to establish that social norms are rational, in the sense of being efficient means to achieve individual or social welfare, or that it is rational to conform to norms, thus reducing compliance to utility maximization.

The first reductionist strategy makes a typical post hoc, ergo propter hoc fallacy, since the mere presence of a social norm does not justify inferring that it is there to accomplish some social function. Besides, it does not account for the fact that many social norms are inefficient, as in the case of discriminatory norms against women and blacks, or are so rigid as to prevent the fine-tuning that would be necessary to successfully accommodate new cases. Even if a norm is a means to achieve a social end, such as cooperation, retribution, or fairness, usually it is not the sole means. Many social norms are underdetermined with respect to the collective objectives they may serve, nor can they be ordered according to a criterion of greater or lesser efficiency in meeting these goals. Such an ordering would be feasible only if it were possible to show that one norm among others is the best means to attain a given social objective. Often, though, the objectives themselves are defined by means of some norm.

Consider as an example norms of revenge; until not long ago, a Sicilian man who "dishonored" another man's daughter or sister had to make amends for the wrong by marrying the woman or pay for his rashness with his own life. The objective was to restore the family's lost

* I would like to thank Russell Hardin for his helpful advice. For their valuable comments, I would also like to thank Jon Elster, Jack Knight, Robert Sugden, and Michael Woodford.

honor, but the social norms dictating the ways in which this could be done were the only means available to identify honor in those circumstances. One may think that some form of monetary compensation would have worked equally well, if not better, in the case in which a marriage was impossible. It would have spared one, perhaps many, lives. But accepting a monetary compensation was not revenge, and since nobody would have ever accepted such an atonement, nobody would have even thought of offering it. Approving of the man who exacts revenge, calling him a “man of honor,” does not necessarily involve approval of the norm as rational or efficient. Even if one thinks a norm unjust and useless, it may be difficult not to conform, since violation involves a collective action problem: nobody wants to be the first to risk social disapproval by breaking the norm openly. That is why people will often break a norm in private but still pay lip service to it in public.

The second reductionist strategy argues that—provided that conformity to a norm attracts approval and transgression, disapproval—conforming is the rational thing to do, since nobody willfully attracts discredit and punishment. If others’ approval and disapproval act as external sanctions, we again have a cost-benefit argument. When there is nobody around to watch what we do and we still conform to the norm—the argument goes—it is because we have internalized these positive and negative sanctions. Yet to maintain that we conform to social norms because of the disapproval involved in violating them is of little help in explaining why norms are there, how they emerged, and why they persist. To say that one conforms because of the negative sanctions involved in nonconformity does not distinguish norm-abiding behavior from an obsession, in which one feels an inner constraint to repeat the same action in order to quiet some “bad” thought, or from an entrenched habit that cannot be shed without great uneasiness. Nor does it distinguish norms from hypothetical imperatives enforced by sanctions, such as the rule that prohibits smoking in public areas. In all these cases, avoidance of the sanctions involved in transgression constitutes a decisive reason to conform, independently of what others do.

The line of argument I wish to pursue is biased in favor of a different kind of reduction. Making norms rational, or making it unconditionally rational to conform to a norm, takes norms for granted. Asking why social norms are there, or why we tend to conform to them, does not shed any light on the norm formation process, since *how* norms emerge is a different story from *why* they emerge and become stable. The thesis I wish to sustain is that social norms are the outcome of learning in a strategic interaction context; hence, they are a function of individual choices and, ultimately, of individual preferences and beliefs. The view that norms are reducible to the preferences and beliefs of those who follow them is not new. David Lewis and Edna Ullmann-Margalit have proposed a game-theoretic account of norms and conventions according

to which a norm is broadly defined as an equilibrium.¹ The conventional game-theoretic account has serious limits, though. For one, it is a static description of norms as clusters of self-fulfilling expectations; it cannot explain, nor was it meant to, how expectations come to be self-fulfilling. The equilibrium account of norms must be supplemented with a story of how interacting agents learn to recognize a behavioral pattern, how they settle upon a stable pattern, and what sort of behavior is more likely to be sustainable.

Learning a behavioral pattern must not be confused with socialization, the process through which the newcomer comes to accept an established group's norm. Since our subject is the development of new norms in a group, learning cannot be separated from emergence. The size of the group matters, though. If learning can easily occur in two-person interactions, it may be impossible in a large population, where all that one observes is aggregate behavior. Norms may emerge through learning in a small group and subsequently spread to a large population by some other mechanism. In the last section of the article, I propose an evolutionary account of the propagation of norms from small groups to large populations.

Finally, an analysis of emergence, as opposed to one stressing the functions fulfilled by social norms, may shed light upon the differences between social norms and other types of regularities, such as hypothetical imperatives, moral codes, or legal norms as well as upon those characteristics which are common to all social norms, however different they might be.

EQUILIBRIUM

Norms serve to guide an individual's behavior but also to allow an individual to anticipate others' behavior. We expect people to conform to norms and expect others to expect us to conform, too. A social norm is, in a way, a cluster of expectations. Expectations, I want to argue, play a crucial role in sustaining a norm. Indeed, conformity to a social norm is not unconditional; it is, rather, a conditional choice based on expectations about other people's behavior and beliefs. One's interest in avoiding the negative consequences of transgression, as well as the feelings of shame and guilt that may accompany it, reinforce one's tendency to conform. But they are not the sole, nor the ultimate, determinants of conformity. Reducing conformity to unconditional utility maximization overlooks the conditional element which characterizes norm-abiding behavior. Besides, approval and disapproval are sanctions that presuppose the existence of norms that everyone expects to be followed. Consider a community that abides by a strict norm of truth telling. A foreigner that, upon entering

1. D. Lewis, *Convention* (Cambridge, Mass.: Harvard University Press, 1969); E. Ullmann-Margalit, *The Emergence of Norms* (Oxford: Oxford University Press, 1977).

the community, systematically violates this norm will be met with hostility, if not utterly excluded from the group. But suppose a large group of liars makes its way into this small society. Probably the truth-telling norm would cease to exist, since the strength of a norm lies in its being followed by almost all of the participants.

It may seem that most people's experience of conformity to a norm is beyond rational calculation. Compliance may look like a habit, thoughtless and automatic, or it may be driven by feelings of anxiety at the thought of what would happen if one transgresses the norm. Yet conformity to a norm may be rational and may be explained in terms of one's beliefs and desires, even though one does not conform out of a rational calculation. As David Lewis himself pointed out in his analysis of habits, a habit may be under an agent's rational control in the sense that should that habit ever cease to serve the agent's desires according to his beliefs, it would at once be overridden and abandoned.² Similarly, an explanation in terms of norms does not compete with one in terms of expectations and preferences, since a norm persists precisely because of certain expectations and preferences: if I ever wanted to be different, or if I expected others to do something different, I would probably overcome the force of the norm.

One is not constantly aware of one's preferences and desires, which are better described as dispositions to act in a certain way in certain circumstances. What is required is that such motives be ready to manifest themselves in the relevant circumstances. If somebody were to ask you now if you prefer a Caribbean holiday and five thousand dollars to a punch in the nose and ten thousand dollars, I do not know what your answer would be. Whatever option you would choose, it is likely that you would never have thought of it before; you would not know, for example, that you preferred the Caribbean holiday and five thousand dollars until you were put in the condition to choose. Analogously, when conforming to a norm, one may be unaware of the expectations and preferences that underlie one's behavior and which become manifest only when they happen to be unfulfilled.

What sort of preferences and expectations underlie the conditional choice to conform? A norm is there because everyone expects everyone else to conform, and everyone knows he is expected to conform, too, but expectations alone cannot motivate a choice. If my compliance is grounded on the expectation of almost universal compliance, it must be that I prefer to comply with the norm on condition that almost everyone else complies, too. When going to a dinner party, I do not wear sneakers, not simply because I expect everybody to wear proper shoes but because I also prefer to wear proper shoes if everybody else does. Note that I do

2. D. Lewis, "Languages and Language," in *Minnesota Studies in the Philosophy of Science*, ed. K. Gunderson (Minneapolis: University of Minnesota Press, 1975), vol. 3, p. 25. See also "Convention: Reply to Jamieson," *Canadian Journal of Philosophy* 6 (1975): 113–20.

not need to assume that the other guests at the dinner party also have conditional preferences. My belief that they will wear proper shoes may be grounded on the idea that they actively dislike sneakers or that perhaps they are very traditional and not given to casual dressing. Of course, I might fear that disappointing their expectations will bring about contempt and thus have some additional good reason to wear proper shoes. But this is an independent, secondary reason.

More generally, a social norm (N) in a population (P) can be defined as a function of the beliefs and preferences of the members of P if the following conditions hold:

1. Almost every member of P prefers to conform to N on the condition that almost everyone else conforms, too.³
2. Almost every member of P believes that almost every other member of P conforms to N.

Conditions 1 and 2 define a social norm as sustained by the beliefs and preferences of those who conform to it; they tell us that a social norm is an equilibrium in the game-theoretic sense of being a combination of strategies, one for each individual, such that each individual's strategy is a best reply to the others' strategies, were the others' strategies taken as given. Each maximizes his expected utility by conforming, on the condition that nearly everybody else conforms to the norm. Note that conditional preference indicates that conformity is not a dominant strategy; if it were, then one would have a reason to conform independently of what other people were expected to do, in which case the equilibrium would not be called a social norm.

A norm is an equilibrium that is supported by a configuration of self-fulfilling expectations: if almost everybody expects most members of P to conform, then almost everybody will conform, given conditional preference for conformity. Jon Elster has argued that the distinction between rational, outcome-oriented behavior and behavior guided by social norms can be upheld by comparing the expectations that accompany norms with those characterizing strategic interaction. In the latter case, "expectations can be derived endogenously from the assumption of rational actors," whereas the expectations involved in social norms are given prior to the interaction.⁴

There are relatively few cases in which one can infer what actors will do from the assumption that they are rational. One such case is that of zero-sum games, in which a player's optimal choice against a rational opponent is a maximin mixed strategy. Another case is a game in which one or more players have a dominant strategy. In both situations ex-

3. For a social norm to exist, it is not necessary to have universal conformity. On this point, see Lewis, *Convention*, p. 97.

4. See J. Elster, "Rationality and Social Norms" (University of Chicago, 1987, mimeographed), pp. 15-16.

expectations can be endogenously derived, provided it is common knowledge that the players are rational. My view is that neither zero-sum games nor games in which there are dominant strategies are suitable models for the emergence of norms.

It is unlikely that norms would emerge out of situations of pure conflict. As Ullmann-Margalit has persuasively argued, even when social norms contribute to maintain a status quo which discriminates in favor of one of the parties, some bargaining occurs. Interests are not diametrically opposed in that the coercing party needs the other's cooperation to achieve his aims: he can threaten a sanction or promise a benefit and is liable to be punished if he does not stick to his part of the bargain.⁵

The case of dominant strategies is quite different. Here at least one player's choice is independent of any expectation he may have, which makes his preference for a given course of action unconditional. Consider the following game:

Game 1

		Player 2	
		Left	Right
Player 1	Top	2, 2	1, 3
	Bottom	1, 5	0, 1

Here player 1 has a dominant strategy, which is to play Top. Note that 1 does not need to have any expectation about player 2's choice, since by playing Top he is always better off, whatever 2 does. His preference for Top is unconditional, as is his rational choice of it. Player 2 instead has no dominant strategy; hence, he has to guess 1's choice in order to make his. Knowing (or believing) that 1 is rational, he can predict that Top will be chosen; hence, he chooses to play Right. Player 2's preference for Right is conditional upon 1's choosing Top, as his choice of Right is rational only with respect to his expectation of 1's play. In this case, player 2's expectations can be endogenously derived by simply assuming player 1's rationality.

Since social norms involve conditional preferences, the presence of dominant strategies would violate condition 1. One would conform irrespective of what others do, but then it would become impossible to distinguish a norm from a habit or a moral imperative. We must conclude that an equilibrium characterization of norms must always take expectations as given, as the conditions under which they can be endogenously inferred from a rationality assumption do not obtain. Here is a different example:

5. Ullmann-Margalit, *The Emergence of Norms*, chap. 10.

Game 2

		Player 2	
		L	R
Player 1	T	2, 1	0, 0
	B	0, 0	1, 2

The game has two pure-strategy equilibria, Top Left and Bottom Right, and each is preferred to the other by one of the players. Because there are no dominant strategies, the assumption that players are rational does not suffice to make a choice. Each player has to guess what the other will do, but what each will do is a function of what he believes the other will do, leading to an infinite regress of expectations. Player 1 prefers Top in case player 2 plays Left, but if he expects 2 to play Right, it is better for him to play Bottom. Similar reasoning applies to player 2. Each has conditional preferences and makes a conditional choice, which is rational only insofar as it is consistent with the expectations he holds.

Suppose game 2 represents the well-known telephone game: two people are talking over the phone but the connection is bad and the conversation gets interrupted. There are two possible ways to continue the communication: either the one who called first calls again or they take turns in calling. If player 1 called first, taking turns is best for him (2, 1), whereas his calling again is best for player 2 (1, 2). Each equilibrium represents a different norm, upheld by a different set of expectations. If the two players happen to share the same expectations, they also share the same norm.

It may turn out that one equilibrium is more conspicuous than the other to the players. For example, one of the two parties may be much older than the other, and the society in which they live might have great regard for seniority. Then it would be tacitly assumed that the younger party calls first. One can imagine many other criteria, each of which identifies an equilibrium as a focal point in Schelling's sense.⁶ But even assuming that no conflict arises over which criterion of choice is appropriate to the case, the existence of focal points begs the question. For a focal point to exist, it must be common knowledge among the players that they describe the game in the same way, but unless it is explicitly assumed, there is no reason to believe that common knowledge exists. If, instead, common knowledge is assumed, the focal point equilibrium as well as the expectations that support it are exogenous to the game; but then a significant part of the coordination problem has been assumed away. What stands in need of explanation is precisely how and whether common knowledge can be achieved and focal points can become such.

6. See T. Schelling, *The Strategy of Conflict* (Cambridge, Mass.: Harvard University Press, 1960).

We have specified the contingencies such that they favor Elster's claim that social norms involve exogenous expectations but not his claim that this is the characteristic that differentiates norm-guided behavior from strategic behavior. Conditions 1 and 2 reconcile norms with strategic behavior, but at a price: characterizing social norms as equilibria spells out the conditions under which norms can be upheld but does not indicate how these conditions can be realized. Since social norms are standards of behavior which have come to be expected by a community in a particular social setting, to describe how expectations become self-fulfilling is part of an explanation of how norms emerge.

NORMS OF COOPERATION

The class of norms I wish to discuss is that of norms of cooperation. These norms play an important role in collective action situations, which are closely related to the n -person Prisoner's Dilemma. In such games, each player has a dominant strategy and rationality dictates choosing it, irrespective of what one expects other players to choose. Specifically, each person can choose whether or not to cooperate, and there is no external authority to enforce sanctions on the defectors. Defection is thus costless, whereas cooperation is costly. Typical examples include voting, polluting, littering, saving electricity during a hot summer, and supporting public causes. In all of these cases, the benefits of cooperation depend upon the number of cooperators. In many of them, this number need not be the totality of those concerned. If enough people vote, refrain from consuming electricity, or support a public cause, all will benefit from the outcome. But those who did not register, kept their air conditioners at full power, or stayed home instead of going around collecting signatures will benefit even more, since they cannot be excluded from enjoying the product of the collective effort of others, while they did not pay any price to start with. In cases such as littering or polluting, where a small number of defectors is enough to do the damage, nearly universal cooperation is needed for the socially desirable outcome to obtain. It is enough that a few people start throwing garbage on a clean beach to induce newcomers to imitate them, since walking to a distant trash can seems futile when the beach is already spoiled. In each case, cooperation involves the risk of a net loss: if too many people defect, those who cooperate pay the cost and reap no benefit.

Rational, self-interested individuals should therefore always defect, even if the collective outcome from joint defection is not Pareto optimal. It would be better for everybody to cooperate, but since cooperation is a dominated strategy, any agreement to cooperate would fail to be self-enforcing, as each player would have an incentive to cheat on the other. Then what stands in need of explanation is the fact that Prisoner's-Dilemma-like situations often do not result in disastrous outcomes; instead we observe that—overall—people tend to cooperate. There are occasions in which cooperative behavior is compatible with rationality. When the

cooperator/defector is not anonymous, cooperation may be expected even in the absence of external sanctions, since it may be in the individual's interest to form a reputation for being a cooperative type. If one lives in a small community, it may turn out that it is better to return the favors one receives, to pay one's debts, and in general to avoid exploiting one's neighbors, since once one has a reputation for being an untrustworthy person, he will never again receive help and will possibly be cheated by the rest of the group. Being cooperative in this case is a form of "global maximization," in that one is prepared to forgo a gain now for a greater future benefit.⁷

Defection should be expected in all those circumstances in which an individual is anonymous, as is the case with large groups such as the firms in a competitive industry or the shareholders of a company. Defection should also occur in small groups, either when it is known that the group will dissolve on a given future date or when some member of the group plans to leave for good. Similarly, if the members of a community believe that it is coming to an end, the belief, however ungrounded, may be self-fulfilling, in that all sorts of defections will be rationally justified.

Interactions that have a well-defined time limit can be represented as finitely repeated Prisoner's Dilemma games in which the players know both the structure of the game and their respective rationalities (in fact, they need to have k -level iterated knowledge of both, where k depends on the length of the game).⁸ In such games, the unique solution is to defect. This conclusion is obtained by working backward from the last play. In the last play, what happened before is strategically irrelevant and there is no future. Hence, the dominant strategy is to defect. At the penultimate stage, the players can infer from what they know about the game and each other what will happen in the last stage. They know that what they choose now will have no effect on their choices at the last stage and therefore choose to defect at the penultimate stage, too. This reasoning leads them to defect at each stage of the finitely repeated game.

When cooperation occurs, it might be due neither to a change in preferences nor to the fact that people commit themselves to nonexploitative behavior. The fact that one's exploitative behavior is likely to be detected and sanctioned by future ostracism is a powerful deterrent, but these interactions have a known time limit. Cooperation becomes less surprising if we think that rationality, far from being a specific, clear-cut mode of action, is an inference to the best choice, given the beliefs

7. See J. Elster, *Ulysses and the Sirens* (Cambridge: Cambridge University Press, 1979), chap. 1.

8. See C. Bicchieri, "Backward Induction without Common Knowledge," in *Proceedings of the Philosophy of Science Association Meeting* (1988), vol. 2, and "Self-Refuting Theories of Strategic Interaction: A Paradox of Common Knowledge," *Erkenntnis* 30 (1989): 69–85. Compare also P. Pettit and R. Sugden, "The Backward Induction Paradox," *Journal of Philosophy* 4 (1989): 1–14.

we have about the circumstances of play. For example, cooperation can result when the players have slight doubts about each other's rationality.⁹

Suppose one player suspects the other to be "irrational" and to play, say, a tit-for-tat strategy, with some small probability $\epsilon > 0$. If the suspected player knows it, she has an interest in confirming the suspicion by avoiding all moves that will reveal that she is rational. Thus she will not respond to defection with cooperation nor will she fail to cooperate following a cooperative move of the other player. Playing "as if" she were a tit-for-tat player, she hopes to induce the opponent to respond "kindly." In this case it is possible to cooperate for a long stretch, the total number of noncooperative plays being bounded above by a number which depends on ϵ . The same result obtains if each player is rational and knows that the other is rational but neither knows that the other knows that he is rational. Then nobody is cheated but everybody has an interest in acting as if he were.

A cooperative equilibrium is supported by a configuration of expectations that makes it rational to choose a cooperative pattern of play. Reason does not favor one course of action over another, irrespective of players' expectations, so if one player expects the other to be less than fully rational, cooperation might ensue. Tit for tat, however, is only one of many possible cooperative patterns of play. For example, in a Prisoner's Dilemma game repeated fifty times, player 1 may decide to cooperate (C) in the first round, and for the next rounds $n = 2, \dots, t < 50$ to choose C in period n unless player 2 chose to defect (D) in period $n - 1$. For rounds $n > t$, he will always defect, regardless of the other player's choice. Were player 2 to play D in period $n - 1$, player 1 will respond with D in period n . He may keep playing D until player 2 chooses C and then play C again. However, he may signal to player 2 his willingness to cooperate by returning to play C immediately after he played D in the previous round. Or they may alternate in playing C and D. In general, since a cooperative pattern is better for both, there will be several cooperative equilibria.

With multiple equilibria, it may be impossible to predict which one will in fact be attained or whether one will be achieved at all by the players. Since cooperative behavior makes sense only under conditions of uncertainty about the other player's type, a cooperative equilibrium presumes that the players make the "right" probability assessments about each other's type. This consideration is particularly relevant since in real life one has to convince the opponent that one is likely to be a tit-for-tat player, while in these games the probability that a player is tit for tat is given and assumed to be common knowledge among the players. Conventional game theory gives no plausible description of how individuals'

9. See D. Kreps, P. Milgrom, J. Roberts, and R. Wilson, "Rational Cooperation in the Repeated Prisoner's Dilemma," *Journal of Economic Theory* 27 (1982): 245-52.

beliefs come to converge, which calls for an understanding of belief formation and evolution.

Usually when we witness "spontaneous" cooperative behavior, we tend to credit it to the existence of norms. Different patterns of cooperation are made to correspond to different norms of cooperation, which can be theoretically represented as focal-point equilibria of Prisoner's-Dilemma-type games. Here, too, the question arises as to how a population playing a given game over and over converges to some stationary equilibrium pattern of play. Unless one takes these norms as primitive, thus falling back into some form of equilibrium reasoning, it remains to be shown how the norms get established and what sort of mechanisms support them.

The equilibrium definition of norms we have provided does not make any distinction between a norm that is followed by relatively few people and a norm that is shared by an entire population. Examples of the first are all those regular patterns of behavior that evolve in families, among friends, or in small, cohesive groups such as clubs and teams. The second type of norm is best illustrated by traffic rules, norms of etiquette, and all forms of racial or sexual discrimination. These latter are norms of cooperation, since they allow a large group, sometimes an entire population, to benefit from excluding some other group from certain activities or goods. What distinguishes the two types of norms is the process through which they come into existence. In both cases, individuals will form some beliefs about other individuals, and if enough individuals share the same beliefs, they will act in a way that will make their beliefs self-fulfilling. In both cases individuals will learn to detect behavioral patterns, but the process of learning will differ according to the size of the group. Even in the absence of communication, in a two-person repeated interaction there is scope for signaling and for experimenting with different actions. In large groups, instead, one's actions go mostly undetected and all that one observes is the aggregate behavior of the group, which is the sole predictor of future outcomes. The individual's influence on the group is marginal, so there is no point in signaling or experimenting.

Although a norm may emerge through learning in a two-person interaction, it may never spread to a population, and if it does, the mechanism accounting for its diffusion is likely to be very different from that which explains its formation. Size will not matter much in those cases in which the passage from few to many individuals does not involve a change of incentives. To illustrate this point, take the case of neighborhood segregation: a white family may prefer to stay in a certain neighborhood as long as other white families stay, so that if everybody expects others to stay, there will be no incentive to leave. When a black family moves in, the immediate neighbors may take it as a sign that further changes in the racial balance will follow. This fear may induce them to move, thereby generating further worries in their immediate neighbors, who

may also decide to relocate. This “snowball” effect is an example of spontaneous coordination: it takes the action of one or two individuals (families, in this case) to generate a collective outcome involving an entire population. The norm “do not live in a racially mixed neighborhood” is an example of a pattern of behavior that, once established among a few individuals, rapidly spreads to larger numbers through a mechanism of self-fulfilling expectations. Moreover, the greater the number of people who move away, the stronger becomes the incentive to move. With norms of cooperation, instead, the incentives to follow the norm are inversely proportional to the number of people involved. This is why it is so difficult to specify a plausible process of norm formation, especially in those cases in which a norm of reciprocity is shared by an entire population.

LEARNING TO COOPERATE

Imagine two individuals engaged in a Prisoner’s-Dilemma-type game which they know will be repeated a finite number of times. They do not know each other nor do they have previous experience with this situation. These people are rational and know that joint cooperation is better than joint defection, but each has no idea of what sort of player her opponent is. After each round of play, each learns how her opponent has played and adapts her subsequent choices to what has been learned. There are many ways a player can adapt, depending on such variables as memory, pattern-recognition capability, and the ability to take into account the effects of her own adjustments upon her opponent’s play.

Let us start by considering the simplest possible case of adaptation: the players are “limited strategists,” in that they simply adapt their choice to the action taken by the opponent on the preceding play. Such agents will not try to identify complex patterns of play nor will they change their strategy in response to another player’s moves.

Each player will start by cooperating/defecting and will subsequently respond with cooperation/defection to the action taken by the opponent. There are eight possible adaptive rules the players may choose:

Rule	First play	Subsequent plays	
		If other did	If other did
		C	D
1.	C	C	C
2.	C	C	D
3.	C	D	C
4.	C	D	D
5.	D	C	C
6.	D	C	D
7.	D	D	C
8.	D	D	D

Rule 1 is 'unconditional cooperation'; rule 2 is 'tit for tat'. Rules 3, 4, 5, and 7 make little sense, as they do not clearly indicate either good will or an exploitative attitude; rather, 4 and 5 seem to indicate a contradictory attitude, whereas 3 and 7 just tell a player to do the contrary of what the other did before. Rule 6 would be adopted by a cautious player, ready to respond cooperatively to a cooperative opponent but unwilling to be exploited even once by a defector; rule 8 is 'unconditional defection', which may be adopted by an overly pessimistic player. We suppose the players' choice of a rule to mirror their psychological propensities, and since we assume the players to be adaptive in a very limited way, we do not expect them to change their strategy in the course of play, since this option would entail far greater learning capabilities on their part. Rules 3, 4, 5, and 7 can be eliminated, since they present combinations of initial moves and adaptive responses that make no sense, and in such a limited adaptive situation they cannot be perceived as "signaling" some complex pattern of play. The players are left with rules 1, 2, 6, and 8 to choose from.¹⁰

Now suppose that the players are playing 100 repetitions of the following Prisoner's Dilemma game:

Game 3

		B	
		C	D
	C	3, 3	0, 5
A	D	5, 0	1, 1

Since the players have four possible patterns of play to choose from, they face the following four-by-four game, in which each of the four rules is a strategy that will be played in the 100 repetitions of the above Prisoner's Dilemma, and the payoffs are the undiscounted sum of the payoffs each player obtains in each repetition of the game:

		B			
		1	2	6	8
	1	300, 300	300, 300	297, 302	0, 500
A	2	300, 300	300, 300	250, 250	99, 104
	6	302, 297	250, 250	100, 100	100, 100
	8	500, 0	104, 99	100, 100	100, 100

This supgame has two equilibria: either both players play a tit-for-tat strategy (rule 2) or they both always defect (rule 8). Note that

10. Brian Skyrms has discussed these rules that he calls "Markov habits" in the context of dynamic deliberation on the part of Bayesian players. If Bayesian dynamic deliberators

such simple adaptive behavior may not allow one to learn the strategy of an opponent; a tit-for-tat player will never learn whether she was matched with an unconditional cooperator or a tit-for-tatter, and an uncompromising defector will never know whether he was matched with a cautious cooperative type. Even if learning were to occur, it would not be exploited to the advantage of the players: a prudent cooperator adopting rule 8 will immediately learn whether her opponent is an unconditional cooperator but will not exploit this knowledge to her advantage, while a prudent cooperator and a tit-for-tat player will “lock” into a punitive pattern, even if, by being lenient just once, either one could induce a dramatic improvement in the overall outcome.

Consider what a player would choose, knowing that she can only marginally adapt. If the opponent is a conditional cooperator, it is better to cooperate, while a tit-for-tat strategy does little harm to a player in case her opponent is the “always defect” type. If matched to a prudent cooperator, it would be better to be unconditionally cooperative, but unconditional cooperation is too risky a prospect, while tit for tat still does better than the remaining rules. Note that tit for tat is a better prospect even if the opponent is “smarter” than the player. Tit for tat will protect the player from exploitation by someone endowed with greater learning capabilities, since his defection will be immediately punished. Hence we would expect adaptive players to choose rule 2 and settle on the 2, 2 equilibrium.

This result, however, crucially depends on the payoffs associated with the outcomes. Consider the following Prisoner’s Dilemma matrix, where the letters represent the payoffs obtained by the row player for each combination of his and the opponent’s strategies:

	C	D
C	<i>a</i>	<i>b</i>
D	<i>c</i>	<i>d</i>

Since it is a Prisoner’s Dilemma, $c > a > d > b$. For n plays of the game, dn is what a defector will score if matched with another defector, while a tit-for-tat player matched with a defector gets $dn - 1 + b$. If $dn - (dn - 1 + b)$ is small, a player will be willing to cooperate, but as $dn - 1 + b$ decreases, the probability of choosing defection increases. If unilateral cooperation were associated with a large loss, the 8, 8 outcome would be more likely to obtain. Note that a player need not assign a high probability to being matched with a cooperator in order to choose a tit-for-tat strategy. In fact, in 100 repetitions of game 3, it is sufficient to assign probability 0.047 to the other player’s being a tit-for-tatter in order

have adequate common knowledge for deliberation, the greater the number of iterations, the greater the degree of mutual distrust needed to justify the selection of the ‘always defect’ rule. See B. Skyrms, “Deliberational Dynamics” (1989, typescript), chap. 6.

to choose that strategy. However, the greater the loss associated with unilateral cooperation, the higher the probability of being matched with a cooperator must be in order to make one choose to be a conditional cooperator.¹¹

Suppose now that the rewards associated with joint cooperation and the punishments associated with joint defection are large enough to justify a cooperative choice in a two-player repeated game, but that the game being played is an n -person Prisoner's Dilemma. If n is large enough to guarantee a player's anonymity, the incentives change and universal defection is to be expected. If, instead, the number of players is such that one's defection can be easily detected and punished, the cooperative equilibrium remains a possibility. The choice to be a conditional cooperator will still depend upon whether unilateral cooperation is not too costly, and if the payoff structure is favorable enough, individual cooperation can be expected to continue as long as the group's past aggregate behavior is cooperative. A behavioral regularity thus established can be highly unstable, though; it will be very sensitive to variations in the payoff structure. For example, since the players' adaptive rule only considers what happened in the previous play, the fact that everyone has "conformed" to a cooperative pattern for a long time has no effect upon the choice of a rule the next time a Prisoner's Dilemma situation occurs; with a payoff structure unfavorable to cooperation, an individual would choose to defect. This conclusion is in line with what has been observed by other writers in the field. Russell Hardin, for example, has pointed out that often a cooperative outcome in the context of a repeated Prisoner's Dilemma is due to the existence of "extraplay" incentives that influence current choices. Such incentives, however, require a certain degree of sophistication on the part of the players.¹² An example of an extraplay incentive is the prospect of future activities involving the same participants, or different ones that will be informed about the past behavior of the players. Under these circumstances, it will be in the individual's interest to create a reputation for being a trustworthy, cooperative type. But reputation effects require that a player perceives his and others' choices as contingent and is able to evaluate the future consequences of present behavior.

What we need is a more complex adaptive dynamics, in which strategic uncertainty and the acknowledgment of the possible effects of one's adjustments on other players' adjustments play a greater role. Consider again two rational players engaged in a Prisoner's Dilemma with a known number of repetitions who do not have any information about each

11. This probability will remain quite small, though. For example, if the payoffs of game 3 are slightly modified so that the loss associated with unconditional cooperation is -5 , the probability of being matched with a tit-for-tatter must be at least 0.069 for a player to choose to play tit for tat in 100 repetitions of that game.

12. R. Hardin, *Collective Action* (Baltimore: Johns Hopkins University Press, 1982), p. 164.

other's type. The players may use introspective analysis to form their prior probabilities about the other's type and in subsequent play try to put those different hypotheses to test.

For example, the players may start by cooperating and see what happens. After a few repetitions, each will have formed some hypothesis about her opponent's strategy and will adjust her strategy accordingly. In a simple, two-person interaction, it is likely that a player will take into account the effects of her own adjustments upon her opponent's play, whereas in a large population, one will take the current state of the rest of the population as a prediction of its state at the next stage, since the effects of one's adjustments on other people's future adjustments are insignificantly small.

In order to eliminate some hypotheses, the players may "experiment" with small deviations, which may or may not be profitable in terms of payoffs.¹³ Consider, for example, the case in which both players cooperate for n repetitions. Then a player may want to "test" the hypothesis that her opponent is a retaliator by defecting in the next round. Her deviation reveals to the other player that she is not an unconditional cooperator, but it might also be taken as a signal that she is unwilling to cooperate from now on, if the number of repetitions is small and the end is not too far away. Testing a hypothesis involves deviating from one's strategy, but it involves the risk of being misperceived by another player as being part of a different strategy. Taking this possibility into account, a rational player will be more likely to experiment at the beginning of the game, in order to restrict the set of possible conjectures about her opponent's type without "confusing" him too much.

Thus a tit-for-tat player may want to ascertain that he is not playing with an unconditional cooperator, and a defector may want to test the willingness of his opponent to "forgive" as well as the severity of his retaliation policy. A prudent cooperator may want to learn if she is playing with another prudent cooperator instead of a defector, whereas an individual that believes he "deserves more" and thus plans to defect, say, twice every three plays, will want to know whether he is playing with a similarly convinced player or with a tit-for-tat retaliator. In the latter case, the two patterns of play may look identical:

Player 1: C D D C D D C D D C D D . . .
 Player 2: D D C D D C D D C D D C . . .

Player 1 may be a tit-for-tat player, but he may also be a conditional cooperator that, like player 2, believes he is more deserving than his opponent. If player 2 initially defects, it will be impossible for him to know whether or not player 1's pattern of play is independent of his choices. A better strategy for player 2 would be to cooperate initially,

13. See D. Fudenberg and D. Kreps, "A Theory of Learning, Experimentation, and Equilibrium in Games" (1988, typescript).

since by the fifth play he will know whether he is playing against a retaliator, in which case he would do better by modifying his strategy and choosing to cooperate until the penultimate play. To see why it is so, consider the following pattern of play of player 2 (who believes he deserves more) against a tit-for-tat player:

Player 2: C D D C D D . . .
 Player 1: C C D D C D . . .

The first three moves of player 1 may suggest a pattern of limited cooperation that tells a player to defect once every three plays, and the fourth move may or may not indicate retaliation, but by the fifth move player 2 will come to see that player 1's next move is always identical to his own previous move, signaling a retaliator that is quick to "forgive" defection and reward cooperation. Note that "tougher" retaliatory strategies will not do as well as tit for tat. For example, a tougher retaliator may delay rewarding the opponent for cooperative behavior if there has been a previous defection by restoring cooperation only after the opponent has unsuccessfully cooperated once. The pattern of play may look like this:

Player 1: C D D D D D C . . .
 Player 2: D D C D D C . . .

Suppose player 2 is an exploitative type who would always defect in the presence of unconditional cooperators unless she is convinced that she is playing with a conditional cooperator who will punish her defection and reward her cooperation, in which case she would maximize her expected utility by cooperating. Since player 1 chooses to cooperate in the first play, and defect in the second, player 2 suspects he is a tit-for-tat player; in order to test her hypothesis, player 2 will deviate from her strategy in the third play, since a tit-for-tat player will respond positively to her signal in the next play. If player 1 keeps defecting, player 2 will know for sure that he is not a tit-for-tat player, but the set of possible strategies he may be playing is still very large. Player 1 may be a conditional cooperator who punishes defection by defecting forever after; he may be willing to signal a cooperative attitude at fixed intervals; he may be a defector who made a mistake in his first move; he may be a "tough" retaliator who will exploit twice or more in return for each exploitative episode he had to suffer; and so on. Depending on the projected costs of undergoing further testing, which will depend both on the assessed probabilities of each hypothesis and the magnitude of the loss associated with unreciprocated cooperation, player 2 may or may not attempt further testing. This example suggests that a "tough" retaliator risks locking himself in a self-defeating pattern, since punishment, to be effective, must be easy to understand.¹⁴

14. I suspect that, whenever an unfair pattern of cooperation emerges, this is more likely to be due to an underlying bargaining game in which one of the parties has greater bargaining power rather than to be the outcome of poor learning.

Prudent cooperation is a difficult strategy to detect, too, since the prudent cooperator will initiate the game by defecting and will subsequently cooperate in response to a cooperative move of his opponent. When two prudent cooperators are matched, their strategies will be indistinguishable from an 'always defect' strategy, and they will keep defecting unless one of them is willing to risk being twice "exploited" in order to test the hypothesis that the opponent is not a defector.

In general, tit for tat has a big advantage over other strategies: it is easy to learn, since it has a clearly recognizable pattern, and it protects the player who adopts it from excessive exploitation by a defector; tit for tat will at best tie, and at worst it will lose no more than one play. Robert Axelrod believes that these features of tit for tat are responsible for its overwhelming success in the computer tournaments he ran.¹⁵ When different strategies were paired off for round-robin tournaments of an iterated Prisoner's Dilemma, Axelrod found out that—if the probability of the game's continuing is sufficiently great—tit for tat scored better than all the other strategies it competed with. What I am suggesting here is that the very characteristics that make tit for tat successful in computer tournaments are also likely to play an important role in all those Prisoner's-Dilemma-like circumstances in which there is repeated interaction but the player is uncertain as to his opponent's character. Especially when the number of repetitions is small and experimenting more costly, tit for tat seems likely to be chosen by a player who wants to signal unambiguously his intentions and benefit from the possibility of joint cooperation.

FROM TWO TO MANY

Once a cooperative equilibrium is established, we may expect it to persist, since data from past experience can be used to predict how an opponent will act in the future. If we learn that we are playing with a tit-for-tat opponent, we recognize that unilateral defection is going to be punished immediately. Under these circumstances, each player will prefer to cooperate if the other cooperates, and each will attach high probability to the opponent playing his part in the equilibrium; hence, each player will have a decisive reason to stick to cooperative behavior. Note that common knowledge has not been assumed nor is it needed to maintain conformity. Since the players will have probabilistic beliefs "close" to the equilibrium but not full knowledge, beliefs will be quasi-consistent, but not necessarily fully consistent, with each other. However, once players' expectations are close to the cooperative equilibrium, they will tend to persist because of reinforcing feedback: each player will play her part in the equilibrium, which will lead the other to expect with greater certainty that the equilibrium accurately predicts what the opponent will do in the future.

Such a stable equilibrium is a norm of cooperation, since it fulfills the conditions that define a social norm as a function of the preferences

15. See R. Axelrod, *The Evolution of Cooperation* (New York: Basic, 1984), p. 53.

and beliefs of the members of the population in which that behavioral regularity exists. It is easy to see how approval and disapproval play only a secondary role in eliciting conformity. If another player defects, one is made worse off in two respects: there is an immediate loss and one is forced to punish the defector in the next round. The obvious disapproval that accompanies defection reinforces a cooperative attitude; it cannot, however, substitute for conditional preferences and beliefs in eliciting cooperation.

We may expect a norm of cooperation to emerge as a stationary equilibrium in a group of players in which the identities of the players and the experiences they have had with each other matter. Once a norm of cooperation has been established in a dyadic interaction, it will tend to persist and elicit conformity in new situations in which both cooperative and competitive strategies are possible. If the subjects involved are the same, or if they carry with them reputations from past play, mutual expectations are likely to be quasi-consistent, since each individual will tend to believe that what has happened in the past is a good predictor of what will happen in the future.

The larger the population becomes, however, the more individuals will tend to ignore the effects of their adjustments on the future course of play of other individuals, as their identity (and reputation) will matter less and less. In fact, I doubt that learning is possible in large, anonymous groups. In large groups, an individual's choice has an insignificant impact on the collective outcome, and defection is likely to go undetected. In those circumstances, experimenting with small deviations from one's strategy makes no sense, since no response is likely to follow. The only data available to predict the future state of the population is its past and current aggregate behavior, so if cooperation has taken place in the past, individuals will tend to expect it to occur in the future, too. In these circumstances, expectations of cooperative behavior will be self-defeating and expectations of defecting behavior self-fulfilling.

If people can learn to cooperate only in dyadic or small-group interactions, the explanation of how norms of cooperation emerge as equilibrium patterns of behavior does not extend to large, anonymous groups, where the presence of conforming behavior might be rather explained by the diffusion of small-group norms through an evolutionary process. Russell Hardin has pointed to the overlapping nature of group activities and the tendency to generalize to similar cases as possible mechanisms through which conventions involving large populations are built up out of dyadic interactions.¹⁶ Examples are the norms of truth telling and promise keeping. One will presumably learn that it is better to be sincere and trustworthy in the context of repeated interactions with the same small group of people and will later adopt the same behavior in situations that are sufficiently similar to the original ones or that involve reputation

16. *Ibid.*, p. 196.

effects, in that violation of the norm is taken to signal a flawed character, and this in turn will jeopardize future interactions. Another striking example of widespread cooperative behavior is the development of the "honor code" that governed international commerce in the thirteenth century. It was common for merchants to buy on credit and clear their debts at some future time; all the seller got was a "bill of exchange," a written promise to pay a sum of money at some future date. Henry Pirenne noted that, since metallic money was scarce, the massive development of commerce was made possible by the practice of credit, which involved the use of bills of exchange as money.¹⁷ It was not the enforcement of a government but rather the trustworthiness of the issuer that backed a bill of exchange and made it usable as a means of payment. All this would have been infeasible had the merchants not been in continual relations of debt and credit with one another and thus concerned with good standing among their peers. One can imagine the original development of norms of business conduct among a few local merchants, their spread to a larger business community through the repeated contacts with foreign merchants provided by international fairs, and finally the emergence of a general, unspoken code of behavior regulating the activities of an international community of merchants and bankers.

An evolutionary model of the spread of a behavioral pattern over an entire population is not in conflict with an explanation of its emergence in terms of individual learning in repeated small-group interactions. Voting, contributing to public charities, and refraining from littering or polluting are choices that are not easily amenable to a rational explanation. They need not be, however, thought of as counterexamples to rational choice theory. They may result from compliance with norms of cooperation that emerged out of rational behavior in other contexts and were subsequently extended to the entire population through selection pressures. The advantages of supplementing a rational choice explanation with an evolutionary approach are twofold. On the one hand, an evolutionary model does not require sophisticated reasoning and learning in circumstances, such as large-group interactions, in which it would be unrealistic to assume them. We may, rather, suppose that some behavioral patterns borne out of strategic interactions spread and evolve in a large population out of simple adaptive mechanisms. It is not too farfetched to assume that strategies that make a person do better than others will be retained, while strategies that lead to failure will be abandoned. Another plausible mechanism is imitation: those who do best are observed by others who subsequently emulate their behavior.

Whether a behavioral pattern that has emerged in a small group will survive in a larger population is an important question to address, and an evolutionary model provides a description of the conditions under

17. H. Pirenne, *Economic and Social History of Medieval Europe* (New York: Routledge & Kegan Paul, 1936).

which social norms may spread. One may think of several environments to start with. A population can be represented as entirely homogeneous, in the sense that everybody is adopting the same type of behavior, or heterogeneous to various degrees. In the former case, it is important to know whether the commonly adopted behavior is stable against mutations. For example, experiments conducted by Axelrod and Maynard Smith and Price show that tit for tat is an *evolutionarily stable strategy*, which means that a population of individuals that adopt it cannot be successfully invaded by isolated mutants, since the mutants will be at a disadvantage with respect to reproductive success.¹⁸ It is also well known that a population of defectors cannot be invaded by isolated cooperators. A more interesting case, and one relevant to a study of the reproduction of norms of cooperation, is that of a population in which several competing strategies are present at any given time. What we want to know is whether the present strategy frequencies are stable or if there is a tendency for one strategy to become dominant over time.

What follows is a simple example of an evolutionary process. A game is repeated n times, and after each round of the game, the actual payoffs and strategies of the players become public knowledge; on the basis of this information, each player is allowed to adjust his strategy for the next round. More formally, let p_{it} be the frequency of strategy i in population P at time t , and let Π_{ij} be the payoff to adopting strategy i if the opponent plays strategy j . Let $\Pi_{it} = \sum_j \Pi_{ij} p_{jt}$ be the total payoff of playing i at time t , which may also be interpreted as i 's fitness at time t . Note that the total payoff is the weighted sum of the different payoffs one obtains by being matched with different types of strategies, where the weights represent the frequencies of those strategies in the entire population at time t . p_{it+1} (p_{1t}, p_{2t}, \dots) represents the frequency of strategy i at time $t + 1$ as a function of the relative frequencies at time t of all the available strategies (including i), and thus depends on the total payoff of playing strategy i at time t , since it is the payoff one obtains by playing a given strategy that determines whether one is going to play it again or to abandon it. The dynamics of the frequency distributions of the strategies can be represented as follows:

$$p_{it+1} = \frac{f(\Pi_{it})p_{it}}{\sum_j f(\Pi_{jt})p_{jt}}, \quad (1)$$

where $f(\Pi_{it})$ is the reproduction rate of strategy i and is a monotonically increasing function of the total payoff of playing strategy i at time t .

An equilibrium is a frequency vector (p_1, p_2, \dots) that reproduces itself over time. A pure strategy equilibrium is one in which only one

18. See, e.g., J. Maynard Smith and G. Price, "The Logic of Animal Conflict," *Nature* 246 (1973): 15-18.

strategy is played by the entire population, that is, it is a vector where $p_n = 1$, and $p_i = 0$ for all $i \neq n$. Of course, any such vector reproduces itself, but what we want to know is whether it is stable against mutations. Given two strategies i and n , the frequency of i over time will decrease if the reproduction rate of strategy n matched with n is greater than that of strategy i matched with n , that is, $p_{it} > p_{it+1}$ in each period as long as the ratio $\frac{f(\Pi_{in})}{f(\Pi_{nn})} < 1$ for each $i \neq n$.

Given a fixed population, the number of mutants rises whenever the fitness associated with the mutant strategy is higher. Following Maynard Smith, the condition that makes an equilibrium stable over mutations is that, for each $i \neq n$, $\Pi_{in} < \Pi_{nn}$.¹⁹

Consider the simple case in which there are only two possible strategies: tit for tat (T) and defect (D). This simplification helps in understanding the dynamic process that becomes much more complex with a higher number of strategies but does not affect the generality of the analysis. Let p = frequency of strategy T and $1 - p$ = frequency of strategy D. We may rewrite equation (1) as

$$\begin{aligned}
 p_{t+1} &= \frac{f(\Pi_{Ti})p_t}{f(\Pi_{Ti})p_t + f(\Pi_{Di})(1 - p_t)} \\
 &= \frac{f[p_t\Pi_{TT} + (1 - p_t)\Pi_{TD}]p_t}{f[p_t\Pi_{TT} + (1 - p_t)\Pi_{TD}]p_t + f[p_t\Pi_{DT} + (1 - p_t)\Pi_{DD}](1 - p_t)}
 \end{aligned}$$

We want to find the solution \hat{p} to the above equation in the nontrivial case in which \hat{p} is different from 0 and 1. That is, we look for a value of p such that—if the frequency of tit-for-tatters in the population is equal to or greater than it—the dynamics will favor tit-for-tatters, and, at that value, the number of conditional cooperators will be stable. Substituting \hat{p} in the above equation, we obtain

$$\begin{aligned}
 f[\hat{p}\Pi_{TT} + (1 - \hat{p})\Pi_{TD}] &= f[\hat{p}\Pi_{TT} + (1 - \hat{p})\Pi_{TD}]\hat{p} \\
 &\quad + f[\hat{p}\Pi_{DT} + (1 - \hat{p})\Pi_{DD}](1 - \hat{p})
 \end{aligned}$$

and, solving for \hat{p} , we get

$$\hat{p} = \frac{\Pi_{DD} - \Pi_{TD}}{\Pi_{TT} + \Pi_{DD} - \Pi_{DT} - \Pi_{TD}}.$$

Suppose there have been n repetitions of game 3. The payoffs obtained by playing T or D against an opponent who plays, respectively, T or D

19. J. Maynard Smith, *Evolution and the Theory of Games* (Cambridge: Cambridge University Press, 1982).




FIG. 1

are the following: $\Pi_{DD} = n$; $\Pi_{TT} = 2n$; $\Pi_{TD} = n - 1$; $\Pi_{DT} = n + 4$. In this case, assuming that n is at least 4, the minimum value at which a group of tit-for-tatters can survive in a population of defectors is

$$\hat{p} = \frac{n - (n - 1)}{2n + n - (n + 4) - (n - 1)} = \frac{1}{n - 3}$$

If n is greater than 4, this is a fraction less than 1. For values greater than \hat{p} , tit for tat will not just survive but thrive, as more and more players




FIG. 2

will see the advantage of adopting it. However, if the frequency of tit-for-tatters is lower than \hat{p} , cooperative behavior will get less and less frequent in each period and will eventually disappear (see fig. 1). On the other hand, if n is less than 4, no matter how many tit-for-tat mutants enter a population of defectors at once, their frequency will decrease in each period and eventually go to zero (see fig. 2).

This simple example illustrates the general point that norms of cooperation that emerged in a small group may extend to a population through an evolutionary process; if the number of repetitions is sufficiently large, a small proportion of cooperators can take over an entire population.

CONCLUSION

The above scenario encompasses several ideas that have so far been kept apart. The emergence of norms is an example of spontaneous order, a form of coordination that takes place without the intervention of a central authority and does not presuppose previous agreements or common knowledge among the parties. Yet social norms embody some form of rationality and are ultimately reducible to individual preferences and beliefs. The traditional game-theoretic account has the merit of emphasizing the interdependent nature of the preferences and expectations that sustain social norms, but does not explain how expectations are formed. What is presented here is a plausible reconstruction of how norms of cooperation can emerge through learning in small-group, repeated interactions and be subsequently adopted by larger groups of people. In this sense, evolutionary and rational choice explanations complement each other.