Chapter 6 Answers From Evolutionary Game Theory

6.1 The Role of Repetition and Reputation

The experimental results reported in chapter 3 demonstrate that individuals are motivated not only by self-interested maximizing of monetary rewards but also by concerns for others and the presumption that others do care for them; that is, they trust others. This result may induce serious doubts in those who adopt a biological perspective and believe that finally only self-interest in the form of our “selfish genes” matters (Dawkins, 1976). There are three main conditions under which concern for others or cooperation between only self-interested individuals can evolve (Sigmund, 2001). First, cooperation becomes possible if an individual’s goal is specified as the maximization of biological fitness, that is, the number of that individual’s own genes transferred to the next generation. In the case of genetically related individuals advantages for others also favor one’s own genes, thereby making cooperation possible (Hamilton, 1964). Group selection is a second possibility for evolving cooperation. This is because affiliation to a successful group might lead to individual benefits (Wilson & Sober, 1994). A third condition that enables cooperation between genetically unrelated individuals is repeated interactions. Repetition enables cooperation because it offers the opportunity to reciprocate cooperative behavior and punish uncooperative behavior. Trivers (1971) was one of the first who emphasized this role of reciprocity. In contrast, in a world with “one-shot”, anonymous relationships concern for others can barely evolve (Sigmund, 2001).

Repetition, as one condition that enables cooperation, has presumably attracted most of the research that tries to explain cooperation between unrelated, only self-interested individuals (Axelrod, 1984; Lewis, 1969; Schelling, 1960; Sugden, 1986; Trivers, 1971; Ulmann-Margalit, 1977). Some of these explanations are in fact subforms of the repetition argument, for instance, that humans might have evolved a mechanism with which they detect whom to trust in interactions (Cosmides & Tooby, 1992). If we can reliably judge when someone has a reputation for cooperating even in interactions without repetition, then it is in our own interest to trust this person and to reciprocate trust to increase our own good reputation. The reputation argument requires repetition so that reputation can be established (see also Frank, 1987; Güth & Kliemt, 1995; Nowak, Page, & Sigmund, 2000). The repetition argument takes us away from the one-shot world with anonymous relationships and leads us to our real world.

The repetition argument gives a simple explanation of why people trust each other and why they reciprocate trust: First, in many interactions trust and reciprocity provide the possibility of a surplus for the engaged individuals. Individual will hesitate to exploit others, because the exploited individual might break off the profitable interaction. Therefore the individuals will strive for an agreement on how the surplus in the ongoing interaction should be distributed. The crucial aspect is how the surplus is distributed among the individuals, in other words, how individuals coordinate on several allocation possibilities (also called the "equilibrium selection problem," Binmore & Samuelson, 1994).

Even if an interaction is nonrecurring but nonanonymous, trust is promoted and exploitation is prevented by the motivation of establishing a good reputation, which is important for subsequent interactions. This reputation argument has been theoretically elaborated on by Kreps, Milgrom, Roberts, and Wilson (1982, see also Kreps & Robert, 1982) and recently picked up by Wedekind and Milinski (2000), who demonstrate experimentally that individuals in "anonymous" interactions give more to those who had given more to others in the past. In sum, repetition can either directly or indirectly promote trust and reciprocity.

I assume that most of our real-life interactions are ongoing or at least nonanonymous. We might have learned a repertoire of decision strategies that are adapted to different social environments. If individuals encounter a particular situation they will use the strategy that appears most appropriate for the situation and adapt it to the situation. If individuals encounter the novel and unfamiliar situation of a one-shot game in an experiment they might use a strategy for a comparable situation, one that usually involves repetition and nonanonymity. Trust and reciprocity can thereby be explained by the use of strategies that have been adapted to the real world, which is usually characterized by ongoing interaction and nonanonymity.

6.2 The End of Adaptation

If one assumes that people's strategies are adapted to their social environment, what do these strategies look like? In other words, how can we predict the kind of strategies that people are likely to use in a particular social situation? One could argue that such an adaptation process in the end will lead to the same equilibria that are the result of a game-theoretical analysis, which makes the identification of equilibria necessary.

A Nash equilibrium is defined as a combination of strategies that are mutual best replies. A strategy is a best reply if it is optimal (i.e. produces the highest payoff) given the other player's strategy. If both players select mutual best replies, then neither of them has an incentive to deviate from his or her strategy. A game-theoretical analysis builds on two principles. The first principle supposes that people have stable preferences regarding the outcomes of their decisions. The second principle says that people act strategically, such that they take the interdependence of their strategies and the strategies of the other players into consideration and behave accordingly. Although these principles sound simple they impose "perfect rationality" on individuals (Samuelson, 2001) and it is questionable whether people really walk through all eventualities of a strategic situation. Alternatively one could argue that people use some initial strategies when they encounter a particular social interaction, which might be based on great contemplation or only on intuition without conscious deliberation, and find out by experience that some strategies perform better than others. During such a learning process people improve their strategies. Of course such a learning process will depend on many aspects: for instance, on how important the decision is for the individual, how much variance there is between the outcomes of the different decisions options, how often the individual encounters the situation so that experience can be acquired, whether the feedback is clear enough to evaluate strategies, and how much opportunity is provided for imitating strategies from other individuals who are more successful. If, for instance, a decision is unimportant for an individual, the decision consequences are similar, and the individual encounters the situation infrequently, it is unlikely that the individual will improve his behavior. However, if the decision is important, the feedback is clear, and the individual encounters the situation repeatedly and observes how other individuals behave, it becomes likely that after a learning process he will apply good strategies, strategies that are most suitable to reaching his goals. This view of bounded rational decision makers implies that individuals often end up with very good decisions even though they might not walk through the game-theoretical analysis researchers apply to analyze a strategic interaction (for a similar view see Binmore, 1998; Samuelson, 1997; Samuelson, 2001). Although such an adaptation process (i.e. individual learning or social evolutionary process) will presumably improve individuals' strategies it is questionable whether optimal strategies are developed (Dupre, 1987; Selten, 1990, 1991). Alternatively they might often end up with good but not optimal strategies or in Herbert Simon's words, with "satisficing" strategies (Simon, 1956).

The above view of bounded rationality illustrates why a theoretical analysis of a strategic situation is important for developing hypotheses about individuals' decision processes. In economics it is the standard approach to study a strategic situation by its game-theoretical "solutions," even thought the underlying assumption about the individuals' rationality is in many situation not fulfilled. In the field of psychology the approach of studying the decision problem and its solutions to derive hypotheses about the decision process also has a long tradition. Early on Brunswik (1956) and Gibson (1979) promoted the idea that human cognition should be understood as an adaptation to the environment. Marr (1982) suggested the idea that an understanding of a cognitive process must include an analysis on the "computation level" of the task and its solutions being performed by the individual. More recently Anderson (1991) has proposed a "rational analysis" approach that is "based on the assumption that it [human behavior] is optimized somehow to the structure of the environment" (p. 471). Anderson emphasized the importance of such an approach because of two main difficulties of cognitive psychology: an "induction problem" that consists of how mental mechanisms can be inferred from behavior, and an "identification problem" that recognizes that many models of cognition are equivalent in their behavioral predictions (p. 471). If one assumes that behavior is adapted to the individual's environment, then the structure of the environment constrains the possible cognitive mechanisms and specifies what behavior a model has to achieve, thereby helping to solve the induction and identification problems. Supposing that individuals are adapted to their social environment, what kind of strategies are the outcome of such an adaptation process?

6.3 Well-Performing Strategies for Social Interactions

We now look at evolutionary simulations that have been used to study the performance of various decision strategies in social interactions. Simple strategies for social interactions have been extensively studied for the prisoner's dilemma. Although there are some crucial differences between the investment game and the prisoner's dilemma the two games also share some similarities: Both are two-person bargaining games with a finite number of possible decisions (assuming that players do not choose mixed strategies, that are strategies which choose between different strategies with particular probabilities) and both offer a surplus for the two players if they "cooperate." Therefore, and because of the large and excellent previous body of research covering the indefinitely repeated prisoner's dilemma, I will describe well-performing strategies for the prisoner's dilemma that are able to explain how cooperation is developed. For the prisoner's dilemma mutual cooperation becomes possible through repetition. In the indefinitely repeated game, both players forego the highest payoff by exploiting the other player in each period and cooperate to sustain cooperation in the long run, which leads to a higher payoff than ongoing mutual defection.

Studying decision strategies that evolve for a particular strategic situation can basically be done in two ways (Sigmund, 2001). First, the "heterogeneous approach" starts with a population of various strategies and makes them subject to selection. Thereby, it explores strategies that could emerge if people encounter a particular game for the first time and learn strategies over time. The learning process depends on the selection mechanism and on spontaneous new strategies. The selection mechanism reinforces those strategies that outperform others and thereby they are learned by the individual. Spontaneous new strategies, including only small modifications of existing strategies, are necessary for improving the existing strategies. The second, "homogeneous approach" starts with a population of identical strategies and investigates whether this population resists invasions of any other mutant strategies. By this approach, one could test whether a particular strategy performs well enough so that any new strategy is less likely to be learned or evolved.

The heterogeneous approach was followed by Axelrod (1984) with his tournament studies, in which researchers from different fields submitted various strategies. To evaluate the performance of strategies, they played the prisoner's dilemma against each other repeatedly. It turned out that the simple Tit-for-Tat strategy on average outperformed all other strategies. The fact that such a simple strategy outperformed many other strategies is quite astonishing and has initiated a broad range of studies. Some of these studies also highlighted some weaknesses of Tit-for-Tat.

In spite of its high performance it should be noted that two Tit-for-Tat strategies are not the only equilibrium in the indefinitely repeated prisoner's dilemma, and a population of Tit-for-Tat strategies can be invaded by other strategies; hence Tit-for-Tat is evolutionarily not stable (Binmore, 1994). One weakness of Tit-for-Tat is caused by the assumption that a strategy's decisions are chosen deterministically (Nowak & Sigmund, 1993). If players make small mistakes when applying a particular strategy, this can have enormous implications on the strategy's performance (see also Selten, 1983; 1988). If, for instance, two players apply Tit-for-Tat and one player defects accidentally, this leads to a long period of disparate decisions, thereby lowering payoffs for both players until one player makes a second mistake. There are various strategies less vulnerable to mistakes. Pavlov, for instance, cooperates if both players made the same decision in the previous period, otherwise it defects. In this way, its behavior depends not only on the other player's decision, as is the case for Tit-for-Tat, but also on its own decision. As when using Tit-for-Tat, two players that apply Pavlov will continue to cooperate as long as both players cooperate. However, if one player defects (by mistake), both players will defect in the following period, but in the next period both players will continue cooperation. Therefore, a mistake affects only two periods and has little effect on the payoffs. *Generous Tit-for-Tat* is another strategy that is less vulnerable to mistakes. It cooperates in the first period, and in the following periods it cooperates whenever the opponent cooperated in the previous period. It defects with a certain probability if the other player defected in the previous period. Nowak and Sigmund (1994) showed that Generous Tit-for-Tat outperforms Pavlov if the prisoner's dilemma is played sequentially compared to the simultaneous game. Another strategy that is less vulnerable to errors is *Contrite Tit-for-Tat*, for which the notion of a "standing" was introduced (Boyd, 1989; Sugden, 1986). Contrite Tit-for-Tat starts to cooperate in the first period, and in the following period it does what the other player did in the previous period, unless Contrite Tit-for-Tat accidentally defects and thereby loses its good standing. In this case, it cooperates to regain good standing and in the next period, it will again cooperate regardless of the opponent's decision. Thereafter, it does what the other player did in the previous period.

The above studies explain how cooperation can evolve and generates a rather roseate picture of the evolutionary dynamics for the prisoner's dilemma. The described strategies are all "kind" strategies as they start cooperating in the first period and do not initiate exploitation. Unfortunately, this picture is not complete: Apart from these types of cooperative strategies, there are other strategies that are less elegant but which also arise frequently in evolutionary simulations. Another simple strategy, *Always Defect*, defects in all periods regardless of the opponent's decision. It is often seen as the noncooperative benchmark; it outperforms Tit-for-Tat in a direct comparison and, therefore, often breaks down cooperation (on the strengths of Always Defect see Boyd, 1989; Marinoff, 1990; Sober, 1992; Young & Foster, 1991). Starting from a heterogeneous population of randomly designed strategies, Always Defect is often the strategy that prevails in the population at the beginning of an evolutionary process, before more complex, cooperative strategies are able to invade the population. Another simple strategy, *Grim*, cooperates in the first period and continues to cooperate unless the other player defects. If the opponent defects, Grim defects in all following periods regardless of the opponent's decisions and will never return to cooperate. Grim's simplicity, its potential for cooperation, and the little room it provides for exploitation, makes it presumably a frequent visible strategy in evolutionary processes. Linster (1992) showed that Grim is the most successful strategy for the indefinitely repeated prisoner's dilemma if the strategies are restricted in complexity and if the evolutionary process is subject to constant mutations. The described strategies provide an overview of how cooperation can evolve for ongoing social interactions and will function as an anchor for the following evolutionary simulations.