Social Interactions and Ethical Values

Chapter · January 2010
DOI: 10.1057/9780230277236_5

2 authors:

Amitava Krishna Dutt
University of Notre Dame

Charles K Wilber
University of Notre Dame

5
Social Interactions and Ethical Values

Introduction

We argued in the previous chapter that individuals have ethical values which influence how they behave. This chapter examines the implications of that behavior for societal outcomes. The outcomes, obviously, depend on how individuals interact with each other. Therefore the main task of this chapter is to examine how the interaction between individuals results in social outcomes, and what role ethical values have in the nature of these outcomes.

Even if we confine ourselves to economic issues, narrowly defined, such as consumption, production, and trade, there are many ways in which individuals interact. Thus, they may interact by sharing things with each other in a family, by making agreements with each other at meetings of firms and workers, by working with each other within a firm, and by buying and selling with each other in markets. Economists mostly focus on interactions in markets, but they also examine interactions in nonmarket spheres, such as in firms, groups, and families.

A technique that has been especially useful for economists in analyzing interactions in general – market and nonmarket ones – is game theory. In fact, game theory has been applied to a variety of contexts, including discussions of war and conflict (such as arms races, and interactions between governments and terrorists), political strategy, and social interactions. In this chapter we will use simple game theory to understand how individuals interact, what outcomes are produced by such interaction, and how ethical values affect these outcomes.

Section 5.1 discusses some general examples of how economic interactions may be affected by ethical norms and ethical behavior. Section 5.2

January 19, 2010 14:45 MAC/DUTT Page-79 9780230_575950_06_cha05

briefly describes games and the prisoners' dilemma and how games involving this dilemma lead to undesirable outcomes. Section 5.3 examines how ethical values and other considerations can overcome such undesirable outcomes. Section 5.4 discusses some issues connected with the emergence and dynamics of norms using game theory. Section 5.5 makes some concluding observations.

5.1 Interactions and ethical values

We begin our discussion with some examples of interactions and the role ethical values can play in affecting the outcomes.

First, consider interactions between two individuals (or two or more groups of individuals) who can cooperate with each other to produce something together and then divide what they produce among themselves. Examples of such possible cooperation may include a range of activities from actually producing goods and services together with a business partnership, a family through marriage, or a political or military coalition. Assume that individuals can devote their time to productive activity – which increases their joint production – or to activity which increases their likely share of the total output they produce. Thus individuals may spend more time on producing, or, for example, they may spend more time acquiring armaments or doing exercises to become stronger. The individuals cooperate to produce more together than the sum of what each could produce individually. If each individual is interested in the share of the output they get, they are both likely to spend some time producing and some time in ensuring that they get a decent share of the product. If they spend most of their time on getting a large share, instead of producing, the total pie will be too small. If they spend most of their time in production, they may be deprived of their share by a partner who becomes very powerful. They therefore end up getting less than what they would if they spent all their time on joint production. In this situation, if they trusted each other, or if they cared not just for what they obtained, but also what their partner obtained, at least to some extent, they would both end up getting more.

Second, consider interactions through markets, where goods are exchanged at prices determined in the market. Because of the large number of issues raised by examining market interactions explicitly, we will discuss them in the next chapter.

Third, we turn to the interaction between individuals within a firm. Consider a firm which employs workers and produces for the market. We will return to this case when we consider markets, since the employer–employee relation is in part a market relation. However, this relationship is not just a simple market relation. Unlike many markets in which the two parties enter into a transaction and that is the end of their relationship, the relationship here is very likely to be a relatively long-term one – even if it is for a month, but it is often for a year – and one in which the firm is not hiring a certain amount of labor services, but a certain number of explicit or implicit hours of labor. When the labor contract is made the employee knows his or her wage, but has not yet provided the labor services. The employer may not be able to monitor exactly how hard the employee works, especially in more complex jobs, even with supervisors, and the employee has a great deal of leeway about how much effort he or she will put into the job. The degree of effort, in turn, will depend on how the employee believes he or she is being treated by the firm. A symbol of this treatment is the wage paid to him or her, although other conditions of work also count. Akerlof (1982) likens this to a gift exchange: if employees believe that they are being well treated and well paid, they will in return be loyal to the employer and buy into the employer's goals. The main point is that if they believe they are unfairly treated, they will not be loyal and feel no sense of duty to get the job done; they will shirk and only work the least amount that they can get away with, and may even sabotage the production process. If, however, they feel that they are treated more than fairly, they will feel satisfied with their job and proud of working for the employer, and therefore put in a great deal of effort. The result will be that if employers believe that this is the way their employees will respond, they will pay a fair wage and will try to provide a good working environment. Productivity and wages would be higher than if they did not pay a fair wage and provide for good working conditions. Moreover, there are likely to be fewer labor–firm disputes, which will have a positive effect on efficiency.

Issues about fairness and trust are important not only in the relationship between the employer and the employees in a firm, but also in the relationships among employees. If there is an absence of trust between employees within a firm, overall labor productivity may be adversely affected. In some cases, especially, when there are different ethnic groups represented in the work force, there may be a low level of trust among workers in a firm. Firms may try to take advantage of the lack of trust among employees to keep wages low, by discriminating between workers from different ethnic groups in terms of wages paid. Although firms may be able to set employees against each other, and thereby even to raise productivity by threatening favored workers with lower wages and providing the incentive to low-paid workers to raise their wages, the resulting climate may be one which is considered unfair, and which may reduce productivity for the firm overall.

Fourth, consider nonmarket relations between firms which are located in one region and which can potentially share technological knowledge. It has been found that firms in a particular industrial area – sometimes called an industrial district – can benefit by drawing on each other's technical knowledge based on what each learns through experience. This is more likely to happen if there develops trust between firms which allows them to exchange knowledge. Using the concept of social capital, Putnam has argued that participation and trust are mutually dependent: "The theory of social capital presumes that, generally speaking, the more we connect with other people, the more we trust them, and vice versa" (Putnam, 1995, 665). As firms trust each other, they can draw on each other's technological knowledge and speed up technological change. This process has been used as an explanation of the success of industrial districts in Italy (see Pyke, et al., 1992) and elsewhere. The process can be explained in terms of simple game theoretic analysis (see below) in which individuals care about each other. However, in business and other relationships, problems may arise because firms are engaged in competition in some spheres (for instance, in competing for markets and skilled workers), and individuals and firms who form associations can harm society by adversely affecting third parties, as in the case of organizations like the Ku Klux Klan and business cartels.

As a final example, consider the issue of managing common property resources. In the widely discussed problem called the tragedy of the commons, it is argued that if there is open access to resources, individuals will overuse them, and the result will be that all individuals will be worse off. A standard solution to this problem that is often recommended is the privatization of common property resources. Some would consider privatization to be inequitable. But is it always efficient? Many common property resources are what are called local common property resources – such as fish stocks, forests, and grazing lands to which relatively few people who live in an area have access – rather than global common property resources, such as the global environment, which is relevant for the global warming problem. Many economists recommend that for the efficient use of these local common property resources, these resources should also be privatized and be owned by individuals. What this overlooks is the fact that, as discussed by Ostrom (1991) and others, in several parts of the less-developed world communities have established informal systems of cooperation which overcome the commons problem through the development of trust relationships. Introducing private property into these communities can destroy these trust networks and impose costs which can lead to inefficient outcomes, for instance, by increasing costs of guarding individual private property. Game theoretic models and other analyses have been developed to illustrate these issues (see, for instance, Seabright, 1993, 2004).

5.2 Games and the prisoner's dilemma

Game theory is concerned with the analysis of situations in which, usually, individuals interact with other individuals and addresses how individuals behave and what the outcome of such behavior is for all individuals.[1] A game is defined by its *players*, the *actions* open to each player, and the outcomes of each action (by each of the players) for each player, called the player's *payoffs*. Games can be expressed in two forms: strategic or normal form games and extensive games. The latter examine moves made by the players, sequentially and one at a time, and are depicted using tree diagrams which show the payoffs for the players after each move by a player. For our purposes it is more useful to examine normal form games which are depicted by tables showing the payoffs for the players for each action by each player. Games are of two basic types – cooperative and noncooperative games. The former applies to situations in which binding agreements between the players can be made and are fully enforceable. The latter applies to situations when these assumptions cannot be made. Here we will only examine noncooperative games.[2]

A simple game with two players with two actions each is shown in Table 5.1. The two players are row player and the column player (so called here because the strategies of the two are shown in the rows and

Table 5.1 A simple game

		Column	Player
		Cooperate	Defect
Row	Cooperate	3, 3	0, 4
player	Defect	4, 0	1, 1

columns, respectively). Each player has two possible actions, cooperate or defect, which refer to whether they decide to abide by an informal agreement they made to cooperate or they decide not to. The payoffs are shown with the numbers in the four cells. By convention, the first number shows the payoff to the row player and the second to the column player. Thus, if the row player cooperates and the column player defects, the row payer gets a payoff of 0 and the column payer gets a payoff of 4. The payoffs may be amounts of money or goods, or may refer to the utility (utils) each player obtains.

This simple game may be "played" only once, in which case it is called a one-shot game, or it may be repeated, in which case it is called a repeated game.[3] If it is repeated forever, the game is called an infinitely repeated game. In repeated games the sequence of actions for a player can be represented as a *strategy*, which gives a complete characterization of the actions of the player in each move. Examples of strategies include: always cooperate; always defect; toss a coin before every move and cooperate if you get heads and defect if you get tails; and start by cooperating and continuing to cooperate if the other player cooperates, but if the other player defects in any move, defect and continue defecting forever, or until the end of the game (a strategy that is called tit for tat). In the case of repeated games the overall payoff to a player is usually obtained by adding up the payoffs for all moves after converting them to their present value using a time discount factor which reduces the value of future payoffs.[4]

One can look for a solution to a game in several ways. One is to look for equilibrium solutions, the most well-known of which is the Nash equilibrium, named after John Nash (b. 1928), the American Nobel Prize-winning economist and mathematician. The Nash equilibrium is an array of strategies, one for each player, such that no player has an incentive to deviate from that strategy since he or she cannot increase his or her payoff by adopting a different strategy, given the strategy adopted by the others.[5] Notice that this concept of equilibrium involves a notion of rationality since players try to do the best for themselves by maximizing their payoff given what others do. But it is not specific about what the payoffs represent: are they amounts of goods and services, are they levels of utility, and do they involve self-interested behavior in the sense that the players only care about themselves (since their payoffs may reflect what other players receive, a point to which we will return below). This notion of rationality will imply that the players will actually choose the strategies associated with the Nash equilibrium. Sometimes there may be multiple Nash equilibria, in which case additional considerations need to be introduced to attempt to find a unique equilibrium.

In the one-shot game shown in Table 4.1 the Nash equilibrium occurs where both players defect and receive payoffs of 1. Neither player will have an incentive to change his or her strategy given what the other player is doing, since by so doing (moving to cooperate) they will obtain a payoff of 0. Note that the players both receive a payoff of 3 if they both cooperate, and both are better off cooperating than they are by defecting. However, the outcome in which both players cooperate is not a Nash equilibrium, since both players will have an incentive to defect and try to increase their payoff to 5. None of the other two possible outcomes, in which one player cooperates and the other player defects is a Nash equilibrium. For instance, in the outcome in which the row player defects and the column player cooperates, the column player can gain by defecting, since his/her payoff by defecting, that is, 1, is higher than his/her payoff by cooperating, that is, 0, when the row player defects.

What makes this simple game very interesting is that the unique Nash equilibrium, in which they both defect, will be the outcome when both players choose their "rational" strategy, but it leaves them both worse off than they would be if both cooperated. If we evaluate what is socially optimal in terms of Pareto optimality, the Nash equilibrium is not a Pareto optimal state, while the Pareto optimal state (where both cooperate where both players are better off than at the Nash equilibrium) is not a Nash equilibrium. This type of game, in which the individual optimization of "rational" agents leads to an undesirable social outcome (which is not Pareto optimal) is called a Prisoner's dilemma game.[6]

Not all games are Prisoner's dilemma games. For some, the Nash equilibrium may also be a Pareto optimal outcome. But situations analogous to the Prisoner's dilemma game may be quite common in the real world. Cooperation may be good for both players if they both cooperate and therefore they may agree to cooperate. But if one player does cooperate and the other player, who has the temptation to break the agreement, does so, the first player becomes a "sucker" and loses a great deal. If agreements are not enforceable, then both players may try to play it safe and break the agreement, to avoid being a sucker. Suppose we interpret the payoffs as amounts of goods and services or money each player obtains, or their utility level which depends positively on what they get. Then having the objective of maximizing their payoff amounts is what we called self-interest. Why, in general, does self-interest not lead to the socially optimal outcome which is Pareto optimal as it does under certain conditions for the perfectly competitive market outcome?

This happens because each player in the game knows that what they do will affect the other player, and vice versa, and will therefore take this into account in making their decisions about what they do. In other words, the two players interact strategically. In perfectly competitive markets each player is so "small" that they assume that what they do will not affect anyone else, and can therefore ignore the reactions to their actions when the decide what to do. Thus, problems arising from strategic interactions do not arise.[7]

5.3 How ethics and other factor affect outcomes in games

Let us suppose that the prisoner's dilemma problem arises in many cases of economic (and other) interactions between people. If individuals are self-interested, socially suboptimal results may well occur. Under what conditions can we avoid such undesirable outcomes? Economists and other game theorists have discussed a variety of possibilities.

One possibility is that we find a third party which can enforce agreements and contracts. The obvious third party is the government and the legal system. The very existence of such an enforcement mechanism may be able to overcome the problem of breaking agreements, and if they are still broken, the player who has been wronged can sue. While this is clearly a possibility in many cases, for various kinds of interactions the payoffs may be relatively small in comparison to the costs of litigation (and the uncertainty of the outcome of the lawsuit) to make this solution practicable. Moreover, this solution may be inequitable if the chances of winning a lawsuit depend positively on how much money is at the disposal of the litigants.

Another possibility which has received a great deal of attention is that cooperation is more likely to occur if there are repeated interactions in which a prisoner's dilemma-type game is played infinitely many times and if people do not discount future payoffs too heavily. The reason why two players are more likely to cooperate in this case is that each self-interested player will take into account the fact that the other player may punish him or her if they defect and make them into suckers by refusing to cooperate in the future. This refusal will make their payoffs in the future lower than they would be if both parties cooperate. If the discount factor is high (that is, players do not discount future payoffs too heavily because they are impatient), there will be little temptation to make a killing now by breaking an agreement. The psychologist Robert Axelrod has found, using strategies that were sent to him by contestants in response to his announcement, that the strategy that provides the highest payoffs in repeated interactions in a round robin in which every contestant plays every other contestant (in a computer simulation), that the tit-for-tat strategy does best in terms of the total payoffs. Axelrod's argument is open to some technical criticisms which need not detain us here. But it is a powerful argument that suggests repeated interactions may sustain behavior which can promote cooperation from self-interested people who do not have any values or norms which induce them to cooperate in a way that is against their self-interest. It also leads to some practical suggestions on how to encourage this form of cooperation, for instance, by encouraging repeated and frequent interactions between people. However, there are problems that may arise with this form of interaction in the real world. If players interact with each other only a few times, or only once, because they are likely to interact with many others in an increasingly globalized world, then the assumption of an infinitely repeated game will not hold. Moreover, if someone is made a sucker and destroyed economically or even physically, then the game will not be repeated with the same payoffs. Further, sometimes a player may interpret a move by the other player as defection, but actually the move was unavoidable, because external circumstances prevented the player from delivering on his or her promise, resulting in a series of retaliatory defections. There are ways around some of these problems, for instance the problem of not having repeated interaction can be addressed through concern for one's reputation by word of mouth and ratings published on the internet, but these need not always result in more cooperation if reputations and ratings can be falsified or are unfair.

A third possibility is that we address the issue that strategic interactions is one of the sources of the problem and change the situation from one in which there are two (or a few) players to one with many players. In this case each individual player becomes too small relative to the entire playing field to have any effect on the other players and hence elicit a response. This change can be achieved by breaking down barriers to entry and by encouraging entry by other buyers and sellers. An example would be to allow more competition from foreign sellers. While this may work in some cases in removing suboptimal equilibria, it will not always do so, as when there are positive feedbacks in the form of externalities. In such cases, society as a whole may end up with socially nonoptimal outcomes even with people pursuing their self-interest. For instance, if people find it beneficial for them to use QWERTY keyboards on their computers (over an alternative, the DVORAK keyboard) and if the number of people already using it is large (which means that they are more familiar with them and that there are many stores which will repair them), everyone may end up using QWERTY keyboards and the DVORAK keyboard will die out, even if the latter allows people to type more quickly, and society would be better off if everyone used it (see David, 1985).

Although these three solutions, and others, may increase cooperation, ethical values also can improve social outcomes. We discuss three possibilities. First, suppose individuals are ethical in the sense that they decide not to tell a lie (not even in response to someone else lying to them). Then, if they say that they will cooperate, they will cooperate. Thus, even without repeated interactions and without government enforcement, the prisoner's dilemma situation will result in people choosing the cooperative solution. This solution will actually be beneficial for both parties even from their self-interest point of view, even if they get no additional utils from keeping their promises and not lying – something that they constrain themselves never to do. Players here do not tell lies not because they want the socially optimal outcome but because they intrinsically value not telling lies. And the result of not telling lies is instrumentally good by making the society more efficient.

Second, suppose that people care about other people, so that their payoffs reflect not only their self-interest, but also what others get. Suppose they value what others get just like they value what they get, then the game of Table 5.1 will be converted from a prisoner's dilemma game to one in which the Nash equilibrium is socially optimal, as shown in Table 5.2. In this game we add the payoff received by the two players for each pair of actions in Table 5.1 and give them to each. For instance, if both players cooperate, and both players are virtuous in the sense discussed here, both will receive payoffs of 6, that is, 3 from what they get and 3 from what the other does. Clearly, cooperation by both becomes

Table 5.2 A game with altruistic individuals

		Column Player	
		Cooperate	Defect
Row player	Cooperate	6, 6	4, 4
	Defect	4, 4	2, 2

the Nash equilibrium. It is also socially optimal if we use the payoffs shown in this table, and we get it even if one views social optimality in terms of self-interested utilities. Note that "full" altruism in the sense that we care about others as much as we care about ourselves, is not necessary to get this result.

In fact, suppose we weigh other people's self-interested payoff only half as much as we weigh our own, we get the payoffs shown in Table 5.3, and again the Nash equilibrium will be socially optimal in both senses. What if the extent of people's altruism in this sense is different? Suppose that the row player is fully altruistic in our sense, but the column player weighs the row player's self-interested payoff only half as much as his/her own. In that case the payoffs will be as shown in Table 5.4. For instance, if both players cooperate, they both receive self-interested utils of 3, which implies that the row players utility level will be $6 = 3 + 3$ and the column players, $4.5 = 3 + (1/2)3$. But even in this case we can see that the unique Nash equilibrium will involve both players cooperating, and both getting their highest levels of self-interested utility.

Table 5.3 A game with partially altruistic individuals

		Column	Player
		Cooperate	Defect
Row	Cooperate	4.5, 4.5	2, 4
player	Defect	4, 2	1.5, 1.5

Table 5.4 A game with individuals with different levels of altruism

		Column	Player
		Cooperate	Defect
Row	Cooperate	6, 4.5	4, 4
player	Defect	4, 2	2, 1.5

Third, consider the ethical value that one should not be envious. Sometimes prisoner's dilemma situations arise not because of the absolute amounts people receive, but because their payoffs reflect envy in the sense that they care about what they get relative to what others get. Tables 5.5 and 5.6 illustrate this situation. Table 5.5 shows the self-interested utility levels that the two obtain from their actions. It can be seen that the unique Nash equilibrium involves both people cooperating, with each player obtaining a level of self-interested utility of 7; we can think of these as amounts of income. Now suppose that the two players are also interested not just in their income, but also what they receive relative to what the other player receives. Suppose, more specifically, that each player's utility is their own income plus four times their income relative to that of the other player. In other words, we have

$$U_i = Y_i + 4(Y_i/Y_{\text{j}}), \quad \text{for} \quad i \neq j$$

where U_i denotes the utility level of player *I*, Y_i the income level of player *I*, and player *j* is the other player. The amended game with the

Table 5.5 A game without envy

		Column	Player
		Cooperate	Defect
Row	Cooperate	7, 7	3, 6
player	Defect	6, 3	2, 2

Table 5.6 A game without envy

		Column	Player
		Cooperate	Defect
Row	Cooperate	11, 11	5, 14
player	Defect	14, 5	6, 6

payoffs embodying this form of envy or relative concerns is shown in Table 5.6. For instance, when the row player cooperates and the column player defects, the row player receives an income of 3 and the column player gets 6. The row player's utility level is then $3 + 4(3/6) = 5$, and the column player's is $6 + 4(6/3) = 14$.[8] It is easy to see that in this case the unique Nash equilibrium is one in which both players defect. Neither player will wish to move to "cooperate" because their level of utility will fall from 6 to 5. However, the Pareto optimal outcome in which both players are better off than at the Nash equilibrium – in which both cooperate – is not a Nash equilibrium because both player will have an incentive to defect and increase their utility from 11 to 14. The modified game is therefore a prisoner's dilemma game. This example shows that if one introduces concern for relative income into a game one can convert it from one which does not exhibit the prisoner's dilemma problem to one that does. It follows that if people are not envious in the sense discussed here, they may both be better off in terms of income and utility (even if this utility included concern for relative income in the evaluation of the outcome, but not in terms of influencing how people act) even if they obtain no additional utility by being good in the sense of not being envious.

5.4 Games and the emergence of norms and ethical values

Game theory can also be, and has been, used to understand issues concerning the emergence and persistence of norms and ethical values. We commented on some general considerations concerning this in Section 4.5 of the previous chapter. In this section we make some brief comments about how what are called evolutionary games can shed light on the evolution of norms and ethical values.

An evolutionary game is a game in which there are several types of players, for instance, in prisoner's dilemma games, those who always cooperate, those who always defect, or in other games, those who are always hawks (who initiate aggressive behavior and not stop it till they get injured or the other backs down) or always doves (retreat whenever an opponent attacks). These players are therefore not optimizers, but follow fixed strategies. The game is played many times, but it is not a repeated game in the usual sense, because the number of players of each type can change over different iterations of the game. The evolutionary feature of this game is that the number of players of each type changes over time according to the payoffs they receive in a particular round. In a particular round the players are assumed to interact with each other following some rules of interaction – such as those involving random pairings of individuals. Players who follow a strategy with a higher return increase in number while those who do less well find their numbers diminish. In natural evolution this type of dynamics occurs because of fitness to the environment and the transmission of genes through biological reproduction, in economic and social phenomena the dynamic may be the result not only of reproduction through habit formation and family upbringing, but also of conscious choices by individuals who change their strategies to emulate more successful ones.

The outcomes of such models have been explored in two main ways. The first method is by finding evolutionary stable strategies, where a strategy is defined as evolutionarily stable if players using it do better against itself than does any other strategy. More formally, we start with a situation in which a particular strategy is played by all players and then a small number of "mutant" players enter the game playing a different strategy. The initial situation is evolutionarily stable if these other strategies do not do better against the first and thereby progressively reduce the number of players using the initial strategy. A second method is by specifying explicitly the dynamics of population shares to see where the dynamic system leads to. The dynamics of population shares of player types may follow different patterns, with one type of player dominating, or with different types of players coexisting, or with cycles in which the share of one type of player increases and then falls, and the process continues in this way.

Axelrod (1984) has shown that if people adopt strategies which, on average, yield higher payoffs in infinitely repeated interactions in one round, more and more people will adopt the tit-for-tat strategy in subsequent rounds, which will result in everyone cooperating (since, according to this strategy, everyone cooperates till someone else does not) although their strategy is not, in fact, always to cooperate. Moreover, an "invasion" by a small number of players who follow a different strategy would not dislodge this equilibrium, making it stable in the limited sense in which the strategy does no worse when played against itself as it does when played against some other strategy. Axelrod uses this to argue that the tit-for-tat strategy will drive out other strategies like "always defect." In fact, everyone will always cooperate, although they will not actually be playing the strategy "always cooperate" but rather the tit-for-tat strategy. People will not be following the golden rule but rather use a rule of retaliating if others do, but everyone will actually be living by the golden rule.[9] This model seeks to explain how cooperation may evolve in a system with prisoners dilemma games which are infinitely repeated and when people do not discount the future too much, even when they are purely self-interested.

Another example which shows how ethical norms can emerge can be found in an evolutionary game in which people compete for shares in a good which can be divided among them. Assume that there is competition between two players in each round and that if their claims are compatible (that is, their claimed shares add up to a number less than or equal to unity) they get what they claim, and if they are not (their shares exceed unity) they wind up fighting and getting nothing.[10] Suppose players of a particular type follow the strategy of asking for a particular share of the good. To make the model simple, assume further that there are only four types of players who ask for, respectively, 1, 2/3, ½, and 1/3, of the good. Then assume that different types of players meet randomly and play the game against each other. At one step of the game, the entire population of players comprises of the four fixed types, and their interactions lead to some payoffs. For instance, if one who wants 2/3 meets one who wants ½, no one gets anything, and if one who wants ½ meets a player of the same type, both will get payoffs of ½, and if one who wants ½ meets one who wants 1/3, they will each take want they want and the remaining 1/6 will go to waste. Suppose the number of people who follow a strategy which does better in terms of their payoffs increases in each subsequent step of the game. Depending on the details of the dynamics, it can be shown that the dynamics of the system will evolve to one in which everyone wants ½, and that this equilibrium is also evolutionarily stable. Thus, while fairness in division can be justified on abstract theoretical grounds, it can also be shown to evolve using this type of evolutionary game-theoretic model.

5.5 Conclusion

Our discussion of how ethical values can affect outcomes has a number of implications. First, the presence of ethical values, which relate to virtues – such as not being envious and caring for others – and to actions – such as not telling lies – may be able to transform social outcomes which are ethically bad to ones that are ethically good. Second, these ethical values may be virtue-related or deontological at the level of the individual, but the good that occurs for society as a whole may be ethical in a consequentialist sense or even utilitarian sense of bringing about the highest level of utility for the highest number. This illustrates the problem that we encountered in Chapter 3 when we tried to classify ethical and justice theories neatly into separate groups. Third, there is no guarantee that good people and good acts will necessarily improve social outcomes in this sense. For instance, if some people are virtuous or do good acts, while others are not and do not, it is not clear that everyone will be better off in the sense of obtaining higher levels of utility. Moreover, if some people do good by helping others, and that reduces the incentives for the latter to help themselves (something that the former did not foresee), it is possible that everyone will get lower payoffs. But that does not detract from the goodness of these acts or these virtues in a nonconsequentialist sense. Trying to obtain high levels of self-interested utility is not the motivation for being good or doing good in the virtue or deontological sense: if it were, these people would be consequentialist and in fact be a member of the *homo economicus*.

Finally, regarding game theory, the models often used in that approach can be argued to be both too narrow and too general. The games that have been discussed in this section and the previous one are very narrow in terms of the assumptions made about the environment in which the individuals interact. For instance, individuals have a given a number of actions, the actions have certain outcomes in terms of payoffs, the payoffs are known to both individuals and known to be known by them, individuals do not have outside options, and the notion of rationality used is open to question. The games are also too general and nonspecific because they do not specify clearly the actual behavior of individuals in actual institutional settings. While some of the assumptions of simple games can be modified in games with more complicated structures, and game theoretical models can be applied to more specific settings, it is an open question to what extent a particular game captures the essence of situations of strategic interaction between individuals in actual economies and societies. There may be no absolute answer to this question, but there is no doubt that game theory models can formalize many kinds of interactions between individuals and show how norms and ethical values can affect outcomes, sometimes making them better in a precise sense, and therefore provide a powerful tool to analyze the role of ethics in individual interactions. However, there are deeper debates about solution concepts used in games, which may limit the applicability of various game theory models, an issue which is too complex to discuss here.