Chapter 6

Norms: Online Appendix

This chapter incorporates insights from what is called “game theory”. Game theory has a longstanding tradition in economics starting as early in the 1940s (Nash, 1951; Osborne & Rubinstein, 1994; Von Neumann & Morgenstern, 2007), but is becoming used more often in other social science disciplines like sociology and political science (Coleman, 1990; Diekmann, 1985; Gintis, 2007; Swedberg, 2001; Voss, 2001), and it has been particularly proved fruitful for the understanding of the emergence of social norms, and norms more generally (Ullmann-Margalit, 2015 [1977]).

The name suggests that it is a theory, but it is actually not a theory in the sense of a coherent set of propositions and assumptions. Rather, Game Theory (GT) is a formal approach that could help you deepen your understanding about the interplay between individuals and their social context under conditions of social interdependency. Game Theory is a particularly useful tool in case the behavior of individuals does not simply add up. If individuals are acting interdependently, then you may want to use GT. Its basic premise is that complex social processes can be understood by presenting them as a simple game between interdependent actors. And this is particular helpful in understanding why certain types of norms emerge.

The most famous of all games in the literature on game theory, is the so-called Prisoner’s Dilemma (PD). This game is about two persons being imprisoned and accused of a serious crime that they did in fact commit. The prosecutor, however, has only hard evidence for a minor crime that they committed. For this minor crime, the penalty is only a year in prison. In an attempt to get the confession of the prisoners for the serious crime, the prosecutor separates the two prisoners and then offers each of them two options: the prisoner could betray the other (and testify that the other committed the crime), or he could remain silent. The prisoners cannot communicate with each other, so what will they do?

To predict the outcome of the decisions of these two persons, and to evaluate whether these outcomes are somehow desirable to them, game theory offers useful tools. If you want to use GT, you need to make a matrix, and then go through various steps. I will discuss them one-by-one.

Step 1: Who are the players?

First, you need to identify the persons or agents of the game, called *players* in the terminology of game theory. Here we have two: Prisoner A and Prisoner B. Importantly, these two individuals represent not two unique persons, but rather two persons *in general*, i.e. any two persons who are in the same situation as prisoner A and B.

Step 2: What are their choice options?
Second, you have to specify their *choice options*. The general rule in GT is to keep things as simple as you can. This means that you realize that in reality, things are always much more complex than in your simple model. But that the model, nevertheless, captures the essential elements of the situation the players are in. In the PD game, therefore, it makes sense to argue that there are only two options for both players: they can either Confess the crime of the other player (and hence betray him) or remain Silent.

Step 3: What are their payoffs?
Third, you must specify the payoffs for each outcome. These can be thought of as the 'rewards', 'costs', or 'utilities' of the players. In Chapter 6, I used a simple scale going from 0 (very unhappy) to 10 (very happy). Sometimes, the outcomes can be expressed in terms of income, but this is not always so. It really depends on the particular case you study. Generally speaking, you need to come up with some 'currency', some 'metric' such as income, which makes sense, and which allows you to rank the outcomes from (highly) desirable to (highly) undesirable. In the PD case, you can use 'years of imprisonment' as your metric. This will result in a so-called payoff matrix (Table 6A). It shows that if Prisoner A confesses (and thereby betrays B), and B remains silent, then A is released from prison (0 years), and B goes to jail for 20 years. The opposite holds true as well: if A remains silent, but B confesses, then A will be incarcerated for 20 years, whereas B will be free. When both prisoners remain silent, the prosecutor has no evidence for the serious crime, and both go to jail for one year. If both testify against each other, each prisoner will go to jail for 5 years.

Table 6A The Prisoner's Dilemma with payoffs in years of imprisonment

		Prisoner B	
		Remain Silent	**Confess**
Prisoner A	**Remain silent**	1,1	20,0
	Confess	0,20	5,5

Note: entries of each cell of a matrix represent the player's payoffs, in this case years of imprisonment. The first number is the payoff of player A, the second is the payoff of player B.

Step 4: What are their strategies?
The fourth step of game theory, is to assess the strategies of the players in the game. To do so, however, you need to think about the micro-level assumptions you make about human behavior. In the classic PD game, a key micro-level assumption is that of "selfishness," which means that both prisoners have self-centered preferences and only care about themselves going to jail as short as possible. In addition, the players are assumed to be rational. This means that they make choices in accordance with their (self-centered) preferences and that they can think ahead in time, and act strategically. Given these assumptions about human action, which strategy do they follow in this situation? The logic of the game unfolds in the following way. Prisoner A reasons that if he remains silent, but Prisoner B confesses, he will be incarcerated for 20 years. If A

confesses and so does B, he gets only 5 years. Thus, assuming that B will confess, it is better for A to confess. But also if B remains silent, it is better for A to speak out: he goes free instead of 1 year imprisonment. Thus, no matter what B does, it is better for A to confess. Therefore, there is a dominant strategy: a strategy that is favorable to choose irrespective of what the other player will do. In the PD-game, the situation is the same for B, which means that both A and B will confess.

Step 5: What is the outcome of the game?

In the fifth step, you can examine the collective outcome of the game. Because in the PD game, the dominant strategy for both players is to confess, the end result is thus that both A and B confess. In the terminology of Game Theory, the players are in a Nash equilibrium –named after the mathematician John Nash. Specifically, there is an equilibrium when no player has anything to gain by changing only his own strategy. Thus, when A makes the decision to confess, given the decision of B, and he gains nothing from changing his decision while B's decision remains unchanged, and the same is, vice versa true for B, then they are in an equilibrium. This means that both players will confess, and that state of affairs is *stable*. They won't change their decisions. The collective outcome of the situation is therefore that both Prisoners go to jail for 5 years.

After you have taken these steps, you can also inquire whether the actual outcome is also the collectively 'desirable' outcome. As can be easily seen, the equilibrium in this case is by no means the desirable outcome of the game for the players. It would have been much better for the players if they both had remained silent, because in that case both would have gone to jail for 1 year instead of 5 years. The combination of the strategies (Silent, Silent) is the Pareto-optimal outcome, also known as *Pareto-efficiency*. It is the optimal outcome, because there is no other combination of strategies in the game, in which one player improves his outcomes, without reducing another player's outcome. In the current game, therefore, the equilibrium (Confess, Confess) is *not* the optimal outcome (Silent, Silent); in other words: it is sub-optimal.

In more general terms, the PD game is an example of a (non-)cooperation problem, which we have discussed in this chapter. In the terminology of GT, in this type of problem the player has the choice to *Cooperate*, C (in the PD game: not to betray his friend, and hence remain silent), or to *Defect*, D (in the PD game: to confess). What is characteristic of the game is the temptation of individuals to free-ride, to choose for the defect strategy, no matter what the other player does. In Table 6B, this Cooperation Problem is presented in general terms, whereby the payoffs are such that: T (Temptation) > R (Reward) > P (Punishment) > S (Sucker).

Table 6B Cooperation Problem in general

		Player B	
		Cooperate (C)	**Defect (D)**
Player A	**Cooperate (C)**	R,R	S,T
	Defect (D)	T,S	P,P

Nash-equilibrium: (D,D). Pareto-optimal outcome: (C,C)

Whereas the classic Prisoner-Dilemma refers to 2 persons (called a *2-person PD*), in social life such PD-problems more often arise in larger groups (*n-person PD*). The reason for this is that in larger groups, an individual cooperative act will make no or only a marginal difference for the collective outcome. Because such contributions to the collective are too small and individuals suffer the costs from their cooperative behavior, they will free-ride and hence defect will be the dominant strategy (Olson, 2009 [1965]).

With the help of game theory, scholars have analyzed the emergence and effectiveness of social norms. The social norm of 'reciprocity' (Gouldner, 1960), for example, is regarded as a solution to social dilemma problems in groups. An example is: 'if someone helps you, help that person in return when they need it', or 'I'll scratch your back and you scratch mine'. This is what some scholars call *positive* reciprocity, but social norms on *negative* reciprocity – retaliation- are abundant in social life as well: 'if someone hurts you, hit him back', 'pay like with like', 'an eye for an eye', and so forth.

How does this social norm of (positive and negative) reciprocity solve the PD problem? The answer is that if people adhere to this social norm, they will play the game in a different way, such that the temptation to defect is reduced and hence the collective is better off. To see this, let's play the 2-person PD game again, but now having in mind that A and B adhere to the reciprocity norm. Imagine that prisoner A chooses to remain silent, hence cooperates. Now what will prisoner B do? He is tempted to defect, and hence confess. But the social norm of reciprocity prevents him from doing this, and hence he cooperates as well. He knows that if he will defect, his partner will pay him back, and defect in the next round. Thus, when both A and B adhere to the reciprocity norm, they will tend to cooperate rather than to defect, thereby leading to optimal group outcomes.

A major assumption, however, for this social norm of reciprocity to emerge is that A and B interact more than once. They should interact with each other repeatedly, otherwise both A and B can never pay back the other when he defects, and both are tempted to defect. In the language of GT, for the social norm of reciprocity to work, it should be an "iterated game" (i.e. repeated interactions), not a "one-shot" interaction.

And indeed, using game theory, scholars find evidence to suggest that the social norm of reciprocity can overcome problems of human cooperation. In computer tournaments organized by the political scientist Robert Axelrod, computers played iterated PD games, and it appeared that the simple Tit-for-Tat strategy (TFT) was most successful (Axelrod, 2006 [1984]). The TFT strategy clearly bring the social norm of (positive and negative) reciprocity in practice. Specifically, TFT always started cooperative on the first move, and then simply copied what the other player (computer) did. Thus, if TFT starts first with C, and then the other player responds with C, TFT will do C again (positive reciprocity: if the other is nice to you, you should be nice to him as well). But if the other opts for D, TFT will do D, too (negative reciprocity: if someone hurts you, hit him back).

Subsequent game theoretical work has found evidence in support for the idea that the social norm of reciprocity is a major social force to overcome cooperation problems. The idea has been expanded in two significant ways.

First, it has been taken into account that people sometimes *misunderstand* each other's motives and behavior (Nowak & Sigmund, 1993; Nowak, 2006). This means that if someone defects, it does not necessarily mean that this person has bad intentions, and hence that he/she should immediately be retaliated. It could well be the case that, quite reasonably, the person in question did not cooperate this time (for good reasons), and he/she will do so on other occasions. Hence, a more mild version of TFT (called *generous TFT*) works better than the original, harsh TFT version.

Second, it has been realized that the social norm of reciprocity can also be effective in *n-person PD games*. Remember that in 2-person PD games, repeated interactions between two players make reciprocity a powerful force, because what player A and B do, will be remembered when they encounter each other in the next round. Hence, if B defects in round 1, then A can pay him back in round 2 by also defecting, and so forth. Such a shadow of repeated interactions between 2 players makes the norm of reciprocity so powerful. But what about human behavior in groups? Why would the social norm of reciprocity work here? In short, the idea is that of *reputation* or *indirect reciprocity*: by being cooperative in small groups, people obtain a reputation of being helpful, and this has the advantage that, in encounters with other members of the group, people will cooperate as well. Conversely, if someone has a reputation for not being cooperative (i.e., choosing the defect strategy), then others will know (gossip, rumors spread in the group), and will choose the defect strategy when encountering that person (Nowak & Sigmund, 2005).

References

Axelrod, R. M. (2006 [1984]). *The Evolution of Cooperation*. New York, NY: Basic Books.

Coleman, J. S. (1990). *Foundations of Social Theory*. Cambridge, MA: Harvard University Press.

Diekmann, A. (1985). Volunteer's Dilemma. *Journal of Conflict Resolution, 29*(4), 605–610.

Gintis, H. (2007). Unifying the Behavioral Sciences II. *Behavioral and Brain Sciences, 30*(1), 45–53.

Gouldner, A. W. (1960). The Norm of Reciprocity: A Preliminary Statement. *American Sociological Review, 25*(2), 161–178.

Nash, J. (1951). Non-Cooperative Games. *Annals of Mathematics, 54*(2), 286–295.

Nowak, M. (2006). Five Rules for the Evolution of Cooperation. *Science, 314*(5805), 1560–1563.

Nowak, M., & Sigmund, K. (1993). A Strategy of Win-Stay, Lose-Shift that Outperforms Tit-for-Tat in the Prisoner's Dilemma Game. *Nature, 364*(6432), 56–58.

Nowak, M., & Sigmund, K. (2005). Evolution of Indirect Reciprocity. *Nature, 437*(7063), 1291–1298.

Olson, M. (2009 [1965]). *The Logic of Collective Action*. Cambridge, MA: Harvard University Press.

Osborne, M. J., & Rubinstein, A. (1994). *A Course in Game Theory*. Cambridge, MA: MIT press.

Swedberg, R. (2001). Sociology and Game Theory: Contemporary and Historical Perspectives. *Theory and Society, 30*(3), 301–335.

Ullmann-Margalit, E. (2015 [1977]). *The Emergence of Norms*. Oxford, United Kingdom: Oxford University Press.

Von Neumann, J., & Morgenstern, O. (2007). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.

Voss, T. (2001). Game-Theoretical Perspectives on the Emergence of Social Norms. In M. Hechter, & K. Opp (Eds.), *Social Norms* (pp. 105–136). New York, NY: The Russel Sage Foundation.