Microeconomic Theory
Lecture 9: Game Theory

Marcel Fafchamps
Oxford University

Readings: Varian, Chapter 15.
Optional reading:
Baland and Platteau, *Halting Degradation of Natural Resources,* OUP, 1996; available at http://www.fao.org/docrep/X5316E/X5316E00.htm

1. Game Theory

General equilibrium theory was the big theoretical achievement of the 1950's and 1960's. After a while, however, economists became disenchanted with the relatively simplistic forms of strategic interaction allowed for by the general equilibrium theory. They longed for more strategic interactions where individual agents, firms in particular, could interact with each other with more forethought than the myopic price-taker automatons of general equilibrium theory or the unsophisticated Cournot and Bertrand oligopolists.

This interest in strategy led to the discovery, by economists, of a body of work known as game theory. Game theory was not initially developed by economists, and it was first successful among the military for war games. But in the last three decades, game theory has been the major area of theoretical development by economists. It found a most fertile ground among economists interested in industrial organizations. It allowed them to go beyond the Cournot and Bertrand models to investigate much more refined strategic interaction.

This immense body of theoretical work has now reached a point of decreasing returns, and economists have become somewhat disenchanted with game theory itself. The principal reason for dissatisfaction is that its predictions critically depend on the minute details of the games themselves. Another often heard critique is that game theory assumes perfectly rational beings. Applied to the game of chess, game theory implies, for instance, that a single, well known solution to chess exists (e.g., white wins, or a draw), and that every player is capable of computing this solution and of playing accordingly... Yet, in real life, nobody knows what this solution is. Put differently, the predictive power of game theory has been questioned because the assumption of perfect rationality is too strong. Current theoretical developments focus on evolutionary games (played by myopic or bounded rational agents) and on experimental results. An active offshoot of game theory is the budding economic theory of networks.

In spite of current disenchantment, game theory remains an extremely useful tool that has enabled economists to venture into areas that were until then terra incognita. In particular, game theory has provided a powerful way of thinking about institutions, laws, contracts, and corporate governance. For this reason, game theory has been most influential in the field of industrial organizations. Game theory has had a very deep impact on development economics because it provides a way of formalizing institutions. More about this in subsequent lectures and modules.

2. Description of a game

The biggest stumbling block in a rapid introduction to game theory is the precise formalism used in this body of work. Let me try to introduce this formalism as painlessly as possible.

- Definitions: game in strategic form and in extensive form. Examples of a game in strategic (normal) form. The payoff of row is the first one. In a game in strategic form, all players move simultaneously.

  Example: Prisoner's dilemma:

		Column	
		Cooperate	Defect
Row	Cooperate	$3,3$	$0,4$
	Defect	$4,0$	$1,1$

  Examples of games in extensive form: tree structure with sequence of moves. Players move sequentially, but need not always know what the other player has played. All games in strategic form can be put in extensive form.

- Assumption of common knowledge: players know each others' payoff and action set. Payoffs denoted $u_r(r,c)$ and $u_c(r,c)$ where r denotes the action of agent row and c denotes the action of agent column. The action set is the set of all possible actions that each player can take. In the prisoner's dilemma, the action set is $\{cooperate, freeride\}$.

- Definition: mixed strategies and pure strategies: play an action for sure (pure), or randomize between two actions with known probability p_c(mixed).

- Definition: subjective probability distributions: beliefs about what the other player will play π_c.

- Definition: best response: choosing one's own action so as to maximize one's expected utility given one's beliefs about what the other will play, i.e., action is:

  $$p_r^* = \arg\max_{p_r} \sum_c \sum_r p_r \pi_c u_r(r,c)$$

  For instance, suppose that player column believes that player row will defect with a 50% probability. The expected payoff from playing cooperate is $0.5 \times 3 + 0.5 \times 0 = 1.5$. The expected payoff from playing defect is: $0.5 \times 4 + 0.5 \times 1 = 2.5$. Defecting is thus a better strategy in this case.

- Definition: rational expectations: subjective probability distribution coincide with actual choices of other players $p_c = \pi_c$.

- Incomplete information: We will not discuss this here but we examine in detail games of incomplete information when we discuss contract theory – e.g., moral hazard and adverse selection.

3. Equilibrium concepts

3.1. Nash equilibrium

A (non-cooperative) Nash equilibrium consists of probability beliefs (π_r, π_c) over strategies, and probability of choosing strategies (p_r, p_c), such that:

1. the beliefs are correct (rational expectations), i.e., $p_r = \pi_r$ and $p_c = \pi_c$, and
2. each player is playing a best response, i.e.,

$$p_r^* = \arg\max_{p_r} \sum_c \sum_r p_r \pi_c u_r(r, c)$$

A Nash equilibrium may be in pure strategies or mixed strategies. If a pure strategy equilibrium exists, it can formally be found by calculating each agent's best response function and figuring for which action they coincide, as we did above for the prisoner's dilemma for a single set of beliefs. It is easy to verify that, for any set of beliefs, defecting is the best strategy in the prisoner's dilemma game. To see this, note that

$$3\pi + 0(1-\pi) = 3\pi < 4\pi + 1(1-\pi) = 3\pi + 1 \text{ for all } \pi$$

Thus, no matter what beliefs are, defecting is the best strategy. This is true for both players. Consequently, the only set of rational expectation beliefs is to believe that the other player will defect, and the only equilibrium is $\{defect, defect\}$.

If the action set is continuous (e.g., choosing a price or a quantity), then a pure strategy equilibrium is formally found by computing each agents' reaction function (i.e., best response function) and finding their intersection. This typically involves writing down an optimization problem for each agent, computing all the first order conditions, constructing a system with all these first order conditions, and imposing rational expectations. The vector(s) of endogenous variables that satisfies(satisfy) all these equations is a Nash equilibrium. The Cournot and Bertrand solutions to oligopolistic competition are examples of Nash equilibria.

Theorem 1. *Every strategic form game has at least one Nash equilibrium.*

This equilibrium need not be in pure strategies; it may be in mixed strategies. Because mixed strategies are not intuitively appealing, this theorem is of limited practical use. Most models constructed by economists have equilibria in pure strategies – often several of them.

- *Example 1: Prisoner's dilemma (PD game):*

		Column	
		Cooperate	Free ride
Row	Cooperate	$3, 3$	$0, 4$
	Free ride	$4, 0$	$1, 1$

Single Nash pure strategy equilibrium $(Freeride, Freeride)$. Note that this equilibrium is not efficient.

- *Example 2: Battle of the sexes:*

		Girlfriend	
		Movies	Football
Boyfriend	Movies	$3, 6$	$0, 0$
	Football	$0, 0$	$6, 3$

  Two Nash pure strategy equilibria: $(Movies, Movies)$ and $(Football, Football)$. Both equilibria are efficient but have different distributional consequences. The difficulty in this game is how to coordinate on one equilibrium. There is also a mixed strategy equilibrium in which the two players randomize.

- *Example 3: Matching pennies:*

		Jack	
		Heads	Tails
Jill	Heads	$1, -1$	$-1, 1$
	Tails	$-1, 1$	$1, -1$

  There is no pure strategy equilibrium: if Jack plays tails, Jill wants to play tails; but if Jill plays tails Jack wants to play heads. Yet there exists a mixed strategy equilibrium: each player randomizes 50/50, i.e, tosses a coin at the same time. Which is why the game is called matching pennies...

3.2. Applications

I illustrate with two examples how the concepts developed above can be used to throw light of important development (and everyday) issues.

3.2.1. Example 1: 'who cleans the canal?'

Suppose two farmers *Abel* and *Bob* share an irrigation canal. This irrigation canal must be cleaned every year. If the canal is not clean, the farmers can only produce rainfed crops, which yield an income of 3. If the canal is cleaned, the farmers can produce irrigated crops, which yield an income of 10. The cost of cleaning the canal is 4. If both farmers participate in cleaning the canal, the cost is shared equally among them, in which case their net income is $10 - (4/2) = 8$. If one farmer cleans the canal alone, his net income is $10 - 4 = 6$. Payoffs can be summarized as follows:

Example: Game of chicken:

		Abel	
		Clean	Free ride
Bob	Clean	$8, 8$	$6, 10$
	Free ride	$10, 6$	$3, 3$

As it turns out, this game configuration is called a game of chicken. It is different from the Prisoner's Dilemma game because, if Abel does not clean the canal, Bob will find it in his interest to clean the canal on his own. And vice versa. This is called a game of chicken – or waiting game – because it is in the interest of both players to wait for the other to do the work. They cannot wait forever, however, otherwise they will not be able to irrigate. Hence the name: the best strategy is to wait for one player to cave in, to be the chicken.

There are two pure strategy equilibria in this game, $\{clean, freeride\}$ and $\{freeride, clean\}$.[1] Each of them is equally efficient, in the potential Pareto efficiency sense. But they obviously have different distributional consequences. Cooperation is not a Nash equilibrium. Neither is it more efficient – in the potential Pareto sense – than either of the two pure strategy equilibria.

There are many situations that can be understood as games of chicken – e.g., who washes the dishes? who stops at the intersection? who picks the trash from the picnic site? etc.

3.2.2. Example 2: the tragedy of the commons

Imagine two herders *Amy* and *Beatrix* sharing a pasture. They are free to choose the number of animals they put out to graze. Of course, the more animals, the less grass, and the less milk they produce.[2]

To illustrate the resulting equilibrium, I begin with a simplified case in which herders can choose two herd sizes, large and small. When both herders choose a large herd, there is overgrazing and they both get a payoff of 2. When they choose a small herd, there is no overgrazing and they both get a payoff of 4. If one chooses a small herd and the other a large herd, there is partial overgrazing but the herder with the large herd gets at more grass since he has more animals, so gets a higher payoff of 5. The herder with a small herd gets 1 – less than 2 because of partial overgrazing.

The payoffs are summarized as follows:

		Amy	
		Small	Large
Beatrix	Small	$4, 4$	$1, 5$
	Large	$5, 1$	$2, 2$

It is immediately clear that this is a PD game: the game has a single Nash pure strategy equilibrium which is (large, large). As in all PD games, this equilibrium is not efficient since $2+2 < 4+4$. This is called the tragedy of the commons: it is in the joint interest of both herders to coordinate not to overgraze. But in the absence of a contracting or enforcement device, they are unable to credibly commit not to do so. Anticipating this, both overgraze.

The example can be generalized to an arbitrary herd size h. To illustrate this, let each herder's profit be written:

$$\pi_i = r(h_1 + h_2)h_i - ch_i$$

where c is the unit cost of animals and $r(.)$ denotes the unit return to herding. We assume that r is a decreasing function of total herd size $h_1 + h_2$ to reflect the overgrazing effect. Profit maximization for herder 1 can be written:

$$\max_{h_1} r(h_1 + h_2^e)h_1 - ch_1$$

where h_2^e denotes what herder 1 thinks that herder 2 will do. The above optimization problem can be seen as computing the best response of herder 1 to the anticipated choice of herd size h_2^e by herder 2. A symmetrical optimization problem can be written herder 2. From these we obtain two first order conditions:

$$\begin{aligned} r(h_1 + h_2^e) + r'(h_1 + h_2^e)h_1 &= c \\ r(h_1^e + h_2) + r'(h_1^e + h_2)h_2 &= c \end{aligned}$$

[1] There is also a mixed strategy equilibrium in which each players chooses to clean with some probability. This equilibrium is more equitable but less efficient because, with some probability, neither cleans the canal.

[2] Or the less weight animals put on.

If we impose rational expectations $h_1 = h_1^e$ and $h_2 = h_2^e$, the above system can be solved simultaneously to obtain the Nash equilibrium level of h_1 and h_2. It is a Nash equilibrium because it satisfies the requirement that both herder plays her best response to the action of the other.

It is possible to solve each first order condition to obtain the optimal choice of, say, h_1^* as a function of h_2. This is called a *reaction function*. Graphically, the intersection of the reaction functions is the point that satisfies both first order conditions – and thus is the Nash equilibrium.

Note the formal similarity with a Cournot model of oligopolistic competition: just replace $r(.)$ with the inverse demand function $p(.)$. Like in that model, the two producers overproduce – they do not take into account the externality they impose on the other producer. In the herder example the externality is through the production process itself. This is called a Marshallian or technological externality. In the Cournot example, the externality is through the market. It is called a pecuniary externality. By similarity with the oligopoly case, the efficient solution (from the point of view of the producers) is to behave collectively, which means maximizing joint profits – behave like a monopolist. This internalizes the externality and eliminates the tragedy of the commons result.

3.3. Subgame perfect equilibrium

For games in strategic forms, the concept of Nash equilibrium is nice and intuitive because all strategies/actions are decided simultaneously. For games in extensive form, however, it no longer is a completely satisfactory concept. Many equilibrium refinements have thus been added in the hope of identifying a single, consistent and convincing equilibrium concept. I think it is fair to say that this search has failed. But it has provided useful and deep insights into the nature of strategic interaction.

One equilibrium refinement concept that one should know about is that of subgame perfection. This concept is concerned with the possibility of *empty threats*, namely, that agents might threaten to take some actions if other players misbehave. Nothing in the definition of a Nash equilibrium prevents empty threats to be effective in supporting an unlikely equilibrum. This is true even though, if other agents did in fact misbehave, the threats would not be put into action because they are against the player's interest ex post.

- Definition of a proper subgame: in games in extensive form, a subgame is a continuation of a game. It is proper if by itself it forms a properly defined game (illustrate – like the branch of a tree).

- Definition of subgame perfect equilibrium: An equilibrium is subgame perfect if it constitutes an equilibrium in all proper subgames.

In a subgame perfect equilibrium, a player cannot threaten to take an action that, if called upon, he or she would not take. (Give example with game in extensive form.) Put differently, this equilibrium concept rules out empty threats – e.g., I will divorce you if you do not wash the dishes – and keep only credible threats, that is, threats that constitute equilibria in their subgame.

There are many other equilibrium concepts corresponding to more complicated settings. In games of incomplete information, for instance, the dominant equilibrium concept is that of Bayesian equilibrium. We will see examples of such equilibria when we discuss contract theory.

3.3.1. Example 1: trust

To illustrate subgame perfection, consider the following game in extensive form. There are two players A and B. Player A moves first. She has to decide whether to trust B. Player B moves next. He can either cheat or not cheat player A. If player B cheats, player A can decide whether to punish or not punish B. Payoffs are as follows:

$$\begin{aligned}\{trust, cheat, punish\} &= (-c, -10) \\ \{trust, cheat, not\ punish\} &= (0, 100) \\ \{trust, not\ cheat\} &= (50, 50) \\ \{not\ trust\} &= (0, 0)\end{aligned}$$

Player A threatens B that if he cheats he will be punished. Given this, it is in B's interest not to cheat – $50 > -10$. However, should B cheat, it is not in A's ex post interest to punish because it costs $-c$ for A to punish while not punishing yields a payoff of 0, which is larger. The threat to punish is an empty threat. Subgame perfection requires that A's strategy in the punishment subgame be a Nash equilibrium. It is not. Hence the threat to punish is not subgame perfect. Hence punishment is not credible. Hence B always cheat. Hence A optimally decides not to trust B. This results in an inefficient outcome.

3.3.2. Example 2: courts and market exchange

The above example can be generalized to any market exchange. In period 1, a seller can choose to transfers possession of an object a to a buyer. In exchange, the buyer promises to pay b in period 2. The buyer values a more than b and the seller values b more than a so that there are mutual gains from trade. If the seller decides not to sell, both buyer and seller have a payoff of 0.

If b does not pay, the seller can do nothing or he can sue the delinquent buyer by incurring a legal cost c. There are two types of buyers: solvent and insolvent. The buyer knows his type but this is not observable to the seller.

With probability p, the buyer is solvent and can pay. In this case, he is condemned to pay $b + f > b$ to the seller, with $f > c$ – meaning that the penalty clause imposed on the delinquent buyer covers legal costs. With probability $1 - p$ the buyer is insolvent, in which case the seller receives nothing and the buyer keeps the good (e.g., it has already been consumed) but goes to jail, incurring a large negative payoff $-J$ with $a - J < 0$.

What is the equilibrium of this game? We begin by noting that an able buyer always pays since, if he did not, he would end up paying $b + f$ instead of b, which is clearly worse. Hence an able buyer always wants to pay, given the threat of court action.

The real question is: can we deter insolvent buyers from purchasing the good? The seller can threaten to sue if he is not paid. If this threat is credible, it is sufficient to deter insolvent buyers since getting the good but going to jail is worse than doing nothing: $a - J < 0$. In a Nash equilibrium nothing precludes the seller from making this threat and this is sufficient to deter cheating.

But there is something intuitively wrong in this example. Say an insolvent buyer purchases the good and subsequently reveals his type to the seller. Is it in the interest of the seller to sue? If he does not, he gets a payoff of 0; if he does, he gets a payoff of $-c < 0$: he sues but he gets nothing. A rational seller would not, ex post, actually sue a delinquent buyer. Hence the threat of court action is not credible. Hence it cannot deter cheating. Anticipating this, the seller will refuse to sell if p is too low so that the chance of not getting paid overweighs the expected gain from trade.

To summarize, in this example there exist a Nash equilibrium that allows trade but by eliminating empty threats the subgame perfect equilibrium eliminates this equilibrium. This shows that courts are not really helpful in an environment in which most buyers are insolvent or have no assets to forclose upon – as is the case among the poor. This can explain why the poor find it difficult to obtain credit, for instance: there is no point sueing them since they have nothing to take (they are judgement-proof); hence they are not punished if they do not pay; hence it is not in their interest to pay; hence no one in their sane mind will lend them money.

3.4. Rivalry, altruism, and social capital

Economists think of utility as a function of absolute income (or consumption), not of income relative to others. Recent literature, however, suggests that individuals derive some satisfaction from consuming more than others, and vice versa. In a two-player game this can be represented for instance as:

$$U_i = \pi_i - \rho\pi_j$$

where U_i is the total payoff of player i, π_i is her absolute payoff, π_j is the absolute payoff of player j, and $0 \leq \rho \leq 1$ is a rivalry parameter. If $\rho = 0$, player i only cares about his own outcome; if $\rho = 1$ player i only cares about the difference between his payoff and that of the other player.

Assuming rivalrous preferences changes nothing to the outcome of the PD game. But it can turn a chicken game into a PD game. To see how, consider our earlier chicken game:

Example: Game of chicken:

		Abel	
		Clean	Free ride
Bob	Clean	$8, 8$	$6, 10$
	Free ride	$10, 6$	$3, 3$

but this time regard the numbers in the above matrix as describing absolute payoffs, i.e., assuming $\rho = 0$. Now let $\rho = 0.5$. The new game looks like this:

Example: Game of chicken with rivalry $\rho = 0.5$:

		Abel	
		Clean	Free ride
Bob	Clean	$4, 4$	$1, 7$
	Free ride	$7, 1$	$1.5, 1.5$

which now has a PD form. The only Nash equilibrium is $\{freeride, freeride\}$. The reason is that, in the presence of rivalry, players hate being 'suckers': they would rather suffer than see the other player enjoy a payoff much higher than theirs. Rivalry is detrimental to efficiency.

Now suppose the contrary, that is, that people are altruistic and derive positive utility from the welfare of others. We now have:

$$U_i = \pi_i + \alpha\pi_j$$

where α measures the extent of altruism. The presence of altruism can turn PD games into games of chicken or other games where the Nash equilibrium is efficient. To illustrate, take the tragedy of the commons game:

		Amy	
		Small	Large
Beatrix	Small	$4, 4$	$1, 5$
	Large	$5, 1$	$2, 2$

but now assume that $\alpha = 0.5$. The modified game with altruism now has the following payoffs:

		Amy	
		Small	Large
Beatrix	Small	$6, 6$	$3.5, 5.5$
	Large	$5.5, 3.5$	$3, 3$

which has a single Nash equilibrium, the efficient equilibrium $\{small, small\}$. The reason is that altruism incides players to internalize the externalities they cause to others. As a result, they value global efficiency more. This tends to favor efficient equilibria.

Economists generally put little emphasis on altruism (outside small, tightly knit groups such as the family) and until recently have not been fully receptive to the idea of rival preferences. They also tend to think of preferences – and hence payoffs – as stable, i.e., not malleable. In this they differ from other social sciences. Psychologists in particular have long been interested in how people's choices can be influenced by various framing effects (i.e., can be influenced by external agents) and how rivalry affects choices in experimental settings. Psychologists have also demonstrated that many human behaviors appear to have an altruistic bend.

In the context of development, these ideas form part of what non-economists often refer to as social capital. The idea is that success in the provision of local public goods – such as cleaning the irrigation canals and refraining from overgrazing – depends on the presence of altruistic feelings among community members and can unravel due to rivalry.

Many social capital enthusiast equate social capital with how cohesive the community is. The examples above illustrate the thought process underlying their approach. Numerous NGO interventions, for instance, aim to 'educate' the poor and to 'raise awareness' about how their actions affect others. They also put an emphasis on 'capacity building', that is, on providing a local leadership that can mobilize community members to be less rival, more altruistic, and hence solve free riding at the community level. Underlying these efforts is the idea that if a local leader can manipulate local preferences away from rivalry and towards altruism, communities can better organize to self-provide much needed local public goods.

This approach is very different from that of economists, who take preferences as given and have a profound distate for any effort to affect people's preferences – which they see as manipulative, morally reprehensible, and ultimately futile because they do not believe that preferences can be changed durably. The debate will undoubtedly continue. But at least you are now equiped to understand its nature and the underlying assumptions various sides of the debate make.

4. Repeated games

PD games are very frustrating. Is there a way out without resorting to the manipulation of preferences? As it turns out, there is. The idea is to repeat the PD game for an arbitrary length of time. By making future play depend on current play, we can bootstrap our way out of the bad PD equilibrium.

To illustrate the intuition behind this idea, consider the tragedy of the commons original payoff which, without loss of generality, I have rescaled so that the Nash equilibrium payoff is 0:

		Amy	
		Small	Large
Beatrix	Small	2, 2	−1, 3
	Large	3, −1	0, 0

and imagine that this game is repeated an infinite number of time. Further assume that players discount the future at rate r and form the discount factor:

$$\beta \equiv \frac{1}{1+r}$$

Players' payoff is now their total discounted payoff until the end of time:

$$\Pi_i \equiv \sum_{t=0}^{\infty} \beta^t \pi_{it}$$

where π_{it} is the instantaneous payoff of player i at time t. If π_{it} is constant, i.e., $\pi_{it} = v$, we note that:

$$\begin{aligned} \Pi_i &= \sum_{t=0}^{\infty} \beta^t v \\ &= v \sum_{t=0}^{\infty} \beta^t \\ &= \frac{v}{1-\beta} \end{aligned}$$

We want to find a punishment strategy that can enforce cooperative play $\{small, small\}$. Our idea is as follows: if both players play $\{small, small\}$ in period t, they play $\{small, small\}$ in period $t+1$. If one of them, say i, plays 'large' at time t, then j plays 'large' forever after – in which case i also plays 'large'. In other words, if any of the two players deviates from cooperative play, they revert to the Nash equilibrium of the one-shot game. This kind of strategy is sometimes called brave reciprocity because players must take a chance for cooperative behavior to arise. Can this punishment strategy deter deviation from cooperative play?

Consider the incentive faced by player i at time t when both players have played cooperatively until then. Should i deviate? If he does, he can get a higher payoff for one period, but gets 0 forever after. If he does not deviate, he gets the cooperative payoff forever after. Which is larger? Formally we have:

$$\begin{aligned} 3 + \sum_{t=1}^{\infty} \beta^t 0 &\lesseqgtr ?\sum_{t=0}^{\infty} \beta^t 2 \\ 3 &\lesseqgtr ?\frac{2}{1-\beta} \\ 1-\beta &\lesseqgtr ?\frac{2}{3} \\ 1-\beta &\lesseqgtr ?\frac{2}{3} \\ \beta &\lesseqgtr ?\frac{1}{3} \end{aligned}$$

If $\beta > 1/3$, that is, if his discount rate $r < 200\%$, it is in the player's interest to cooperate. What this shows is that cooperation *can* be sustained in communities that will live together for a long time. Note, however, that non-cooperation forever is also a subgame perfect equilibrium here – so are many other more complicated equilibria. Nothing guarantees that cooperation will be achieved, but it is made possible by repetition.

As it turns out, the above results hinges critically on the assumption of infinitely lived players. This may be a bit of wishful thinking... With finitely lived players, the punishment strategy unravels. To see this, consider what happens in the last period. In that period, there is only one equilbrium: non-cooperation. Knowing this, what is the equilbrium in the period before last? Non-cooperation as well. And so on by backward induction.

Fortunately, there is a way to salvage the punishment strategy by assuming that players leave with a constant probabilty $1 - \theta$. We can then redefine $\beta = \frac{\theta}{1+r}$ and the earlier result is recovered. The reason is that players cannot tell which is the last period, and so they expect the game to continue with some probability, and thus may find it in their interest to continue cooperative play.

Repeated game theory has been very influential in development economics, especially in the modelling of market interaction and of mutual insurance arrangements.