Chapter 13: Strategic Choices: An Introduction to Game Theory

I. Introduction to Agent Interdependence

For the most part, the analyses of the first century or so of microeconomics focused on choice settings in which an individual's or organization's decisions fully determined the outcomes of interest. Consumers determine how much of product X they want to purchase, and their decision fully determines how much they do in fact buy. Firms determine what, how, and how much to produce and sell, given market prices or the demand curve confronted, and that is exactly what happens. They produce in the intended things in the intended manner and sell the intended quantities.

An exception to that rule was our analysis of duopoly and related Cournot-types of markets, where the price was neither exogenous, as it is for price takers, nor was it fully controlled by the firm, as it is for price makers. Instead, it was the joint outcome of two firms in the duopoly case and more than two firms in the oligopolistic cases. The same could be said about fully competitive markets, but in the latter case—the price taking cases—the effects of one's purchase or output decisions on prices are so small that they can be neglected without loss.

In cases in which market prices are jointly determined by a relatively small number of firms, the market prices that arise are no longer entirely determined by their own choice or exogenous, but jointly determined by a firm or proprietor's own choice and that of at least one other proprietor or firm. Each firm can be said to make strategic decisions—decisions that depended partly on expectations they have about other individuals or organizations that are not necessarily fully known beforehand.

As the effects of market structure (number and sizes of firms) on market prices attracted the attention of economists, what came to be called conjectural variations—beliefs about how one's rivals would respond to their choices—became an integral part of the theory of price determination in these intermediate forms of markets. By taking such expectations into account, that subarea of economics might be said to have invented the field of game theory—although that would be a stronger claim than could be supported, because historically there have been numerous settings in which such expectations played a role in decision making in noneconomic domains.

For example, the battlefield decisions of generals always depend partly on their expectations about the future choices that would be made by their rivals. However, there were relatively few systematic mathematical treatments of those choice settings. And, thus, the field of industrial organization can be said to have invented game theory as far as economists were concerned.

It was not until World War II that game theory emerged as a separate field of study in applied mathematics. The book that brought the field to the attention of persons outside the small group of applied mathematicians initially working on game theory is the *Theory of Games and Economic Behavior* by von Neumann and Morgenstern (1944). A second more accessible classic work was published a decade later, *Games and Decisions* by Luce and Raiffa (1957). During the next few decades, the game theoretic approach was taken up by a subset of economists, sociologists, political scientists, biologists, and military strategists.

The use of game theory to analyze economic problems continues to be one of the most active areas of contemporary microeconomics. A quick look at any economics journal (and many other social science and philosophy journals) in the past three decades will reveal a large number of articles that use elementary game theory to analyze economic behavior in a variety of market, social, and institutional settings.

Although game theory is a relatively new area of specialization, the use of game theory in economics is older than game theory itself. As noted above, the Cournot's duopoly model (1838) reviewed in chapter 5 is an early example of a non-cooperative game with a Nash equilibrium. Stackelberg's (1934) analysis of duopoly with sequential entry is also game theoretic, as are essentially all models of oligopoly and monopolistic competition. Many other applications have been worked out in the past half century or so.

Contemporary research on self-enforcing contracts, credible commitments, contract renegotiation, externalities, the private production of public goods, and models of political and social activity all have used game theoretic models as "engines of analysis." Most of these applications rely upon the rational choice models (utility maximizing or net-benefit maximizing models) that we have used throughout this text. Game theory thus belongs in any textbook on microeconomics. It can be used to characterize a wide variety of interactions that contribute to the size, scope, and routines of market networks.

As true of neoclassical models, game theory models normally assume that players are "rational" in that the selection of strategies can be modeled as utility maximizing or net benefit maximizing choices. The models developed are instances of methodological individualism in that social outcome emerge from the joint decisions of the individuals in the game, contest, or choice setting of interest.

Formally, games may be one-sided as in the early neoclassical models, but analyses of such choice settings using game theory benefit little from game theoretic terminology, models and results. The games that attract the interests of social scientists are all ones in which the strategy choices of the players in the game jointly determine the outcome and the payoffs (utilities or net benefits) actually received by the players. Game theory can be used to analyze strategies for board games, gambling, or card games—and some of the ideas associated with probabilistic settings were doubtless first worked out in such informal settings—but for the most part, the games of interest are not parlor games, but choice settings in which a good deal of interdependence exists among the participants, and outcomes are jointly determined.

The choice settings modeled using game theory are usually ones in which there are a few "players." This is not because the results depend on small numbers, but because most results from such games can be easily extended to games with large numbers of contestants and repeated versions of the same games. The players themselves may or may not realize that the interdependency exists, but it is an important and consequential aspect of the choice settings of interest to social scientists.

This chapter provides an introduction to game theory. It begins with some elementary but very influential choice settings with just a couple of players and strategies, and then proceeds to settings with more strategies and more players. The illustrating two-person examples are ones that have influenced developments in microeconomic theory and many of its subfields.

The chapter focuses on non-cooperative games in "normal form," all of which can also be characterized using decision trees (as models in "extensive form")—but for students interested in such characterizations, books focused on game theory rather microeconomics provide better sources than can be developed in a single chapter. A few such books are listed in the references. The appendix includes some of the more technical vocabulary of game theory and provides a proof of the existence of Nash equilibria in very general circumstances.

Games consist of players, possible strategies, and payoff functions. Payoff functions characterize the payoffs (usually utilities or profits) realized by each player for all possible combinations of strategy choices. In addition, games specify information sets (often implicitly) that characterize what players know about the game and their rivals.

Most economic applications focus on non-cooperative games in which the players choose their strategies simultaneously, without knowing what the other players will do. Many of the choice settings that economists use game theory to model are presented as "one-shot" contests in which the players are strangers, the game occurs only once and without linkages to other games. However, repeated games are also often of interest, and many have the same equilibria under plausible assumption. (See the appendix for a discussion of sub-game perfect equilibria.)

II. Two-Person Two-Strategy Games

The simplest games that allow one to model social interdependence are two-person two-strategy one-shot games. A surprising number of insights about markets, firms, cooperation, and the nature of competition can be obtained from such simple games. Most graduate students in economics will be familiar with several such games, but we'll nonetheless start out with a bit of review for those students and use the games reviewed as an introduction to game theory for those that have not had a course in game theory.

Two-person two-strategy games can provide useful insights about how self-interested forward-looking individuals will interact with one another in a variety of choice settings. The choice settings are characterized by the payoff functions, the information set, and the strategy set. There are only four possible outcomes to a two-player two-strategy one-shot game:

(1) both players may choose strategy S_1,

(2) both may choose strategy S_2

(3) player A may choose S_1 and player B may choose S_2,

(4) player A may choose S_2 and player B may choose S_1.

The four outcomes can be represented in a "game matrix," a table with four cells, each of which corresponds to one of the four possible outcomes. The entries in the cells are the payoffs that each player realizes, usually in terms of utility, but other payoffs may be used as well such as income, profits, output levels, or net benefits.

In non-cooperative games, each player independently chooses their strategies. The outcome of the game is a consequence of those choices and is represented in a game matrix as the cell associated with the two strategies chosen, which provides the payoffs realized by both players, given the combination of strategies chosen.

The 2x2 game matrices provide a very neat way of characterizing a wide variety of choice settings and problems that individuals in such settings may confront—especially those without internalized norms (a topic taken up in Chapter 16).

A Trading Game

The first game reviewed in this chapter is what might be called the trading game. In that choice setting, player A (Alice) may be regarded as the seller, and she decides whether to offer a product for sale or not. Player B (Bob) may be regarded as the buyer, and he determines whether to accept the seller's offer or not. Initially, we'll assume that if an offer is made and accepted both players are better off. If an offer is not made or not accepted, no trade takes place, and the payoffs are (0,0) for each. No change in their initial well-being occurs. We'll also assume that there are transactions costs in this choice setting. In the setting of interest, it is costly to make offers and also to search for and accept any offers made.

It is the relative size of the payoff numbers that generally determines the strategy choices made by each player. Any series of numbers with the same relative magnitudes would generate the same equilibrium outcomes (assuming that a higher payoff is always more desirable than a lower payoff—as true of utility functions by definition).

Table 13.1: A Trading Game

Bob

		Trade	Don't
Al	Trade	(a, b) (3, 2)	(a, b) (-1, 0)
	Don't	(0, -1)	(0, 0)

Alice has apples and Bob has money. Alice is thinking trading some apples for some of Bob's money. Bob is thinking about trading money for some of Alice's apples. Since trade is voluntary, nothing happens unless both players agree to trade. For the purpose of illustration, it is assumed that it costs "one util" to make an offer, whether it is taken or not and it costs "one util" to try to accept an offer, whether it is made or not. If Al makes an offer and Bob accepts, then their respective payoffs are 3 (for Al) and 2 (for Bob). If Al does not make an offer and Bob does not seek one out, then nothing happens and the payoff for each is zero. If Al makes an offer but Bob does not seek one out, then Al's payoff is -1 and Bob's is zero. If Bob seeks out an offer but Al does not make one, then Bob's payoff is -1 and Al's is zero..

A non-cooperative game is said to be in a **Nash Equilibrium** whenever a strategy combination is "stable" in the sense that no player in the game can realize a higher payoff by changing his or her strategy, given that of the other player or players.

There is not always a unique Nash equilibrium. The above trading game **has two Nash equilibria**, (trade, trade) and (don't, don't). Neither person can make themselves better off by changing their strategy (alone) if he or she finds themselves in one or the other of these two cells of the game matrix.

Which equilibrium emerges is not clear. If Al has chosen not to offer any apples for sale, then Bob should not look for such an offer (0 >-1). If Al has offered apples for sale, then Bob's best strategy is to accept the offers (2>0). There are mutual gains from trade in this choice setting: (3>0) and (2>0). Thus, one equilibrium is better than the other, but that does not alter the fact that there are two stable outcomes in this choice setting. Some markets never emerge. (Think, for example, of all the goods in your closet that could have been listed for sale and sold on *Craigs List, Facebook Marketplace,* or a similar service, but that are not.)

Adam Smith once said that trade emerged in Great Britain during his time partly because individuals and families had propensities to "truck and barter," which is to say partly because many persons enjoyed trading itself—as many persons today enjoy shopping online or strolling through department stores or malls. In such cases, rather than a negative payoff from making offers not taken or seeking offers not made, Alice and Bob would realize a payoff of -1+B, where B is the subjective benefit from shopping. If B is larger than 1, then the mutually beneficial trading outcome would be the unique equilibrium in this form of the trading game.

In such cases, propensities to truck and barter would enlarge the scope of mutually beneficial trade and cause trading networks to be larger than they would otherwise have been.

Normative Assessments of Nash Equilibria

Games are often used to illustrate choice settings that are normatively relevant. When undertaking normative analysis, economists typically use one of two normative theories. They may use the Pareto criteria: An outcome is said to be **Pareto optimal** (or Pareto efficient) if and only if there is no feasible alternative that would make one person better off without making another worse off. The other is the utilitarian norm, which judges outcomes on the basis of the sum (or some other function) of the utilities realized—the higher the better. In a game matrix, that norma can be applied whenever payoffs are in utility levels—if one believes that utility numbers can be added up. In that case, one can use a Benthamite aggregate utility function to assess the outcome that emerges as the Nash equilibrium. An outcome is efficient or optimal from that utilitarian perspective when it maximizes the sum of the utility levels.

Note that the trading equilibrium in game above is both Pareto optimal and utilitarian optimal . No feasible change (e.g. none of the other three cells) make one person better off without making another worse off. In fact, all the other possibilities make both potential traders worse off. Also, the sum of the payoffs (utilities) is maximized at the trading equilibrium.

There are two possible equilibria, but one is normatively more appealing than the other.

The Prisoners' Dilemma Game

The Prisoners' Dilemma game is the most widely used choice setting in social science. This is not because social scientists are fascinated by criminals, but because the problematic incentives created by the payoffs of the classic prisoner's dilemma are present in a wide variety of socially relevant choice settings

The "original" or "classic" prisoners dilemma game goes something like the following. Two individuals (Al and Bob) are arrested under suspicion that they took part in a serious crime (murder or grand theft). Each is known to be guilty of a minor crime (say shoplifting), but it is not possible to convict either of the serious crime unless one or both of them provides evidence about the other's participation in the crime. The prisoners are separated. Each is asked to testify against the other, in which case the person testifying will receive a somewhat lower punishment and the other will be convicted of the more serious crime and so receive a greater punishment than that associated with the minor crime they both were arrested for.

The payoffs are in utility terms with negative numbers indicating a worsening of their condition relative to the status quo that existed before they were arrested. (Some illustrations will use years in jail or the magnitude of fines as payoffs, in which case the players would try to minimize the payoff rather than maximize it.) The payoffs are again listed as (Al's, Bob's).

Table 13.2: The Prisoner's Dilemma

		Prisoner Bob	
		Testify	Don't
Prisoner Al	Testify	(-9 , -9)	(0, -10)
	Don't	(-10, 0)	(-1,-1)

Each prisoner is promised a one-year reduction in penalty if he or she testifies against the other. In this case, there is single Nash equilibrium where they both testify against each other, and both receive a modest reduction in their consequently much larger penalties. Of course, both would have been better off if neither had testified against the other, but the incentives of their choice setting induce each to testify against the other.

To see this, consider the payoffs of Prisoner Al. If Al thinks that Prisoner Bob will not testify, Al has a choice between a strategy the yields a payoff of -1 (not testifying) or 0 (testifying). So, Al is better off testifying. Testifying is Al's "best reply" in that case. On the other hand, if Al thinks that Bob will testify against him or her, then Al has a choice between a strategy that generates -10 (don't) and one the generates -9 (testify). Again, testifying is the best reply.

In this case, Al is said to have a **pure dominant strategy**—no matter what the other person does, he/she is best off with a single unique strategy—namely testify. The Nash equilibria of games in which both players have pure dominant strategies are easy to determine.

The prisoner's dilemma game is usually assumed to be symmetric (the payoffs and strategies are the same for each player for the various combinations of choices), and thus the same logic applies for Prisoner B. Testifying is always the best reply for each, and thus both prisoners decide to testify, and the result is the upper lefthand cell. (This is the intersection of their best-reply functions.)

In this PD-type of choice setting, each pursues their self-interest, but the result is worse for both than if both had cooperated with each other and refused to testify against the other.

This is the Prisoner's dilemma. Although society may be better off (assuming they are both actually guilty of the more severe crime), the prisoners are clearly worse off than if they had cooperated and both kept quiet.

The prisoner's dilemma game (PD game) is widely used in economics and other social sciences because there are a number of settings in which the payoffs (incentives) resemble those of the PD choice setting. Examples include:

i. Competition between Bertrand (price setting) duopolists.

ii. Decisions to engage in externality generating activities. (Pollution)

iii. Competition among students for high grades in universities

iv. Public Good Problems

v. The arms race

vi. The Hobbesian Jungle (Anarchy)

The essence of a PD game's dilemma is that the "cooperate, cooperate" solution is preferred by each player to the "defect, defect" equilibrium; but, nonetheless, the "defect, defect" outcome emerges from independent decision making, because it is the "best" decision for each player no matter what the other player(s) chooses to do.

This contrasts with independent choices in markets, as within the trading game, where independent decision making tends to make the participants better off—or at least no worse off than they were originally.

PD payoffs are normally represented using "ordinal" utility levels with (3, 3) for the mutual cooperative solution and (2, 2) for the mutual defection result (the two "on-diagonal" outcomes). The other payoffs (the off-diagonal outcomes) are often represented as (1,4) and (4,1) with the defector receiving 4 and the cooperator 1. This generates outcomes that are analytically similar to those of the classic contest. In order for the "defect, defect" outcome to be a unique equilibrium, defecting has to yield a payoff that is a bit higher than the off-diagonal payoffs (2>1) and defecting, when the other person cooperates, has to have a larger payoff than realized by mutual cooperation (4>3).[1]

The PD game's main limitations as a model of social dilemmas are its assumptions about the number of players (2), the number of strategies (2), the period of play (1 round).

It is easier to imagine cooperation emerging in small number settings than in large number settings because the choice setting may be repeated, and trust may emerge and the payoffs from a series of such games may be such that the cooperation makes sense.

However, many PD-like choice settings have more than 2 players which makes such cooperation less likely. Many others are also repeated through time, which may allow various patterns of behavior to emerge. (See the Appendix for a brief discussion of the Folk theorem.)

Analytically, the number of players and repetition do not matter as long as defection remains a pure dominant strategy and repetition is finite. In settings in which games are repeated and exit is possible, players that never defect may flourish. (See, for example, Vanberg and Congleton, 1992). The exit option transforms the game matrix from a 2x2 matrix to a 3x3 matrix. (The number of sells depends on the strategy sets available to the players.)

Other "Named 2x2 Games"

Several other 2x2 games have attracted attention and have been widely enough used to have been given names.

i. A **zero-sum** game is a game in which the **sum of the payoffs in each cell** is always zero. In this game, every advantage realized by a player comes at the expense of other players in the game.

ii. **Coordination games** are games where the "diagonal" cells (top left or bottom right) have the same payoffs(for example, 1,1) which are greater than those in the off diagonal cells, (for example, 0,0). Such games illustrate why it can be useful for

---

[1] The PD payoffs can also be represented algebraically using (abstract) payoffs, where (C, C) and (D, D) are the payoffs of the mutual cooperation and mutual defection outcomes, respectively. The payoff of a cooperator who suffers from defection of the other player is often termed S, for "sucker's payoff" with T being the "temptation payoff" for defection. In a PD game, T>C>D>S.

conventions to emerge—e.g. a norm that is followed by both persons. For example, if we all drive on the left side of the road or all on the right, we all have higher payoffs (experience fewer accidents) than if people randomly drive on each side of the road.

iii. **Assurance games** are similar to coordination games. The diagonal payoffs for the "cooperative" strategies are again above those of the off-diagonal cells; however the upper left-hand "cooperative" cell has a higher value for both players (3, 3) than the lower right-hand "uncooperative" cell, which is often regarded to be the original position (2, 2). Note that the trading game can be regarded as a special case of an assurance game.[2]

iv. **Chicken** games are contests in which identical strategy choices are disastrous rather than beneficial. In Chicken games, the "on-diagonal" strategies have lower payoffs than the "off-diagonal" outcomes. The classic chicken game involves two drivers driving in the dark down a country road, with each driver starting in the same lane. The person that changes lane is considered to be a "chicken." If neither chickens out, the cars crash and both players may die or be severely injured. The off-diagonal payoffs for each participant are higher, although one person does better than the other as with (4,2) and (2,4), because he or she has not "chickened out."

III. Economic and Political Economy Applications of PD-like Game Matrices

Game matrix characterizations of two-player games are not limited to symmetric games nor to contests where the strategy sets include only two strategies, nor to ones where the strategy sets are similar. Extensions to three or more strategies are often useful to explore what might

[2] Some game theorists have renamed the assurance game a stag-hunt game after a setting described by Rousseau in his Discourse on Inequality (1755). "If it was a matter of hunting a deer, everyone well realized that he must remain faithfully at his post; but if a hare happened to pass within the reach of one of them, we cannot doubt that he would have gone off in pursuit of it without scruple and, having caught his own prey, he would have cared very little about having caused his companions to lose theirs." See Brian Skyrms' (2004) book on the Stag Hunt for a complete discussion. (Note that the above translation suggests that the Stag-Hunt is really a PD game rather than an assurance game, at least if the starting point is two hunters at their stag hunting post. In the usual assurance game, the better equilibrium is stable, if it emerges.)

be considered intermediate solutions in settings where many strategies are possible. Using game matrices to model such contests often sheds light on the properties of contests with continuous strategy sets without the restrictions that continuous characterizations of payoff functions often impose to facilitate the use of calculus.

The Shirking Dilemma and Team Production

Although production by teams can be highly efficient, there is a sense in which team production is unnatural. Each team member's effort increases the productivity of other team members, but these effects can often be ignored by a person who decides to goof off a bit rather than fully devote him- or herself to team production as characterized by the team's rules. Every person on every team has private incentives to underprovide his or her services to the team. They are inclined to "shirk" rather than "work."

To illustrate this dilemma, suppose that a team is organized as a "natural cooperative" and shares the output produced equally among team members. Each person participates in the team's activities for eight hours. For purposes of illustration, assume that the team output is two times the total effort invested in production (work effort). Suppose, however, that an individual's effort is unobservable to others--such as when a group tries to lift or carry a heavy object, separately searches for fruit to harvest and share, or jointly develops a complex computer or phone app. The benefits of leisure in contrast to work effort (the absence of productive effort) are realized only by the person(s) shirking.

The game matrix illustrates the "shirking" dilemma for a two-member team with three possible strategies. (Two is, of course, the smallest possible team.) The payoffs in the game matrix are net benefits measured in output units. They are the sum of each team member's share of the team's output plus the value of each player's own leisure. The value of an hour of shirking to the individual benefiting from it is assumed to be equivalent to 1.5 units of the team's output. Note that this choice setting has a single Nash equilibrium at the lower right-hand corner of Table 13.3. A good deal of shirking takes place in equilibrium.

Table 13.3: The Shirking Dilemma of Team Production in Natural Cooperatives (with Hours of Effort as Strategies)

		Harold		
		8 hours	6 hours	4 hours
Armen	8 hours	(A, H) 16, 16	(A, H) 14, 17	(A, H) 12, 18
	6 hours	17, 14	15,15	13, 16
	4 hours	18, 12	16, 13	14, 14

That a problem exists is implied by several normative theories. There are other feasible outcomes that would make all team members better off. For example, both Armen and Harold would benefit if they both diligently worked eight hours each day instead of four.

To the extent that shared output or shared profits are correlated with utility levels, aggregate utility is not maximized. And, to the extent that the output of the team contributes to a village's survival by increasing its material reserves, the shirking dilemma diminishes the village's likelihood of survival in the long run.

Solving such intra-firm dilemmas is one of the reasons that firms exist. The shirking dilemma tends to be larger for larger teams, because there are more persons to monitor and coordinate, but the essential problem also exists for small teams—unless internalized norms or external penalties alter the incentives faced by each team member.

Reciprocal External Cost Problems

As noted above, a variety of social dilemma problems have payoff structures that create PD-games like incentives. A choice-setting that is important for the field of public economics is one in which reciprocal externalities exist. In the two-person version of that choice setting, two individuals are permitted by law to engage in an activity that they find profitable or pleasurable, but which indirectly imposes costs on the other. For example, two neighbors may independently host barbeques that produce a good deal of smoke, and play loud music that accords with the tastes of the one playing it, but not their neighbor's. Another similar problem is mowing one's lawn early enough or late enough to interfere with their neighbor's sleep. Other larger scale varieties include congestion on highways and large-scale forms of air pollution.

A 3x3 matrix is more appropriate for external-cost problems than 2x2 matrices because middle solutions are often the ones that are Pareto efficient or maximize the sum of the utilities of the individuals involved, rather than all-or-nothing types of solutions.

The choice setting modeled below can be thought of as one where Bob likes country music and Alice loves opera, or as two separate barbeques, with more or less smoke and odors imposed on their neighbor. The strategies are "full volume," "half volume," and none (no noise or no smoke).

Table 13.4: A Reciprocal External Cost Game

		Bob		
		Full	Half	None
Al	Full	(A, B) 2, 2	(A, B) 4, 1	(A, B) 5, -2
	Half	1, 4	3, 3	4, -1
	None	-2, 5	-1, 4	0, 0

In matrices larger than a 2x2, it is often useful to underline each player's best replies to the strategies that the other may adopt, because there are often several equilibria. Nash Equilibria occur in the cells where both player's payoffs are underlined, which in this case is in the upper lefthand corner. There is just one equilibrium in the game depicted, because each player has a pure dominant strategy, namely, "full." There is unique Nash equilibrium of very loud discordant music or very smokey barbeques.

There is a Pareto superior move from that equilibrium to the cell in the center. (A **Pareto superior move** makes at least one person better off and no one worse off.) The center cell is both Pareto optimal and a Benthamite optimum (it maximizes the sum of the utilities realized). But it is unstable, because each is better off with "full" if the other has adopted "half." Thus, any neighborly agreement might well be broken, and a return to the upper lefthand corner.[3]

The same matrix can be modified slightly to show how rewards and penalties can be used to solve such problems. Assume that the town has adopted a fine of F dollars for excessive noise or smoke which reduces the payoff associated with "full" by f utils. (Dollars and utils are of course quite different things; thus, the use of F and f.) We now incorporate the effective penalty, f, into the game matrix. The fine is borne only by the persons choosing the "full" strategy.

Table 13.5: A Possible Solution to the Reciprocal External Cost Game (with f >1)

		Bob's Possible Strategies		
		Full	Half	None
Al's	Full	(A, B) 2-f, 2-f	(A, B) 4-f, 1	(A, B) 5-f, -2
Possible	Half	<u>1</u>, 4-f	<u>3</u>, <u>3</u>	<u>4</u>, -1
Strategies	None	-2 , 5-f	-1, <u>4</u>	0, 0

Note that a fine that reduces the full payoffs by more than one util is sufficient to solve the problem. How large the fine F must be to do so (e.g. to generate an f >1) is not completely obvious, but it may not have to be very large. Given that effect, the best replies (underlined) are now "half" for both Alice and Bob, and the new Nash equilibrium is the middle (3,3) cell, where the best-reply functions intersect. The parties go on but are not as loud or smokey. (Note also that no fines are actually collected in equilibrium, but the fines are nonetheless important.)

Political Economy Dilemmas

Although firms, clubs, and governments often adopt policies that reduce or solve various dilemmas confronted within those organizations and communities, there are also cases in

---

[3] It should be noted that external benefit problems also exist in which, instead of "over supply" the problem tends to be "under supply." Game matrices can easily be created to illustrate such problems as well, although the external cost problem attracts a good deal more attention than external benefit dilemmas.

which the decisions of one organization or community imposes costs (or benefits) on other organizations or communities. In the next sections, two illustrating examples are developed. (Both have attracted sufficient interest among economists to have been named.)

The "Race to the Bottom." Suppose that there are two communities that are interested in regulating some activity within their own territory. Suppose further that regulations in each community affect each other's prosperity, with the community with the "weakest" regulations being somewhat more prosperous than the community with stronger regulations. To simplify a bit, assume that there are just three types of regulations that can be imposed: weak, medium, and strong regulations. Suppose also that the joint ideal is "medium, medium." However, the economic effect of local regulations (as business relocated to the less regulated community) implies that each community is a bit better off weakening its regulations, given the other's regulation of the activity of interest.

Table 13.6: The Race to the Bottom Dilemma

Community B's Environmental Regulations			
	weak	medium	strong
A's env regs weak	A, B 6, 6	A, B 8, 4	A, B 9, 2
medium	4, 8	7, 7	8, 5
strong	2, 9	5, 8	6, 6

Such games have a Nash Equilibrium in pure strategies that is not necessarily Pareto Efficient. Given the numbers used, the middle equilibrium is better for both, but incentives induce both communities to weaken their regulations beyond that level. Hence, the so-called "Race to the Bottom" dilemma.

However, it is clear that the nature of the contest varies with the payoffs assumed: the (strong, strong) cell could also be an equilibrium, according to the assumptions made about how their joint payoffs vary with the degree of regulation in the neighboring community.

Being able to see both the dilemma and the fact that it is not the choices per se that generate the dilemma, but rather the economic or political relationships behind those payoffs that determine this is one of the advantages of using discrete strategy sets and small numbers of players to illustrate dilemmas. Continuous versions of the same contest would make the key features of the choice settings a bit less obvious.

"Race to the Top," also known as **NIMBY (not in my back yard)**. Now suppose that the inter-community externality in the opposite direction. Suppose that the community with the weaker regulation attracts undesirable (say, noisy, ugly, or polluting) industries. If voters would prefer such industries to locate elsewhere, they would vote in favor of more stringent regulations. We'll again assume that there are just three levels of regulation and that the two-community ideal is (medium, medium) as in the previous example. A very minor change in the payoffs can transform the previous game into the NIMBY game. In this case, each community is just a bit better off if it has somewhat tougher regulations than its neighbor.

Table 13.7: The Race to the Top Dilemma

Community B's Environmental Regulations

	weak	medium	strong
A's env regs	A,B	A,B	A,B
weak	6, 6	4, 8	2, 9
medium	8, 4	7, 7	5, 8
strong	9, 2	8, 5	6, 6

This game also has a Nash Equilibrium with dominant strategies that is not Pareto Optimal. But instead of under regulation, over regulation emerges as the equilibrium in this case. Both are conceptually possibilities.

Notice also that both types of choice settings have clear effects on the production methods being undertaken and the extent of the markets in each community. It is through such effects that the regulatory externality problems arise.

Such problems are political ones in that they are chosen by government officials or by voters via a referendum. Nonetheless, they have economic consequences, and therefore are at least partly within the field of economics.

It is important to note that solutions are not as easy as they look. For example, a voluntary agreement to move to (medium, medium)—as with a treaty—may not solve the dilemma because it is not a Nash equilibrium. Both governments have incentives to violate such agreements. It may be for this reason that international environmental treaties often have little effect on international air pollution [See, for example, papers by Murdoch and Sandler (1997)].

In a federal system of government, the shift of such regulations to a higher level of government might be a solution, but such shifts are generally not feasible in international settings.

IV. Games with Continuous Strategy Options

The above analysis of games with small numbers of players and strategies are analogous to the concrete function form models of Part I. They capture important features of many commonplace choice settings, but are often less general than more abstract models of the same choice settings. In many cases, strategy sets are continuous rather than discrete and in at least a subset of such cases, the continuous nature of the strategy sets is an important aspect of the choice setting.

As true of the more abstract models of Part I, modeling settings in which the domain of strategies is continuous often yields more general and convincing results than possible with strategy sets composed of only a few discrete strategies. Moreover, the comparative statics of those equilibria tend to be less "jumpy" and more easily analyzed than in the discrete cases.

These more general and realistic choice settings begin by characterizing a payoff function that characterizes each player's payoffs (utilities) as a function of the strategy choices of all the players in the game of interest.

We'll begin with two concrete function forms similar to ones used in Part I to illustrate how this can be done.

The Shirking Dilemma

The previous section included a 3x3 game matrix representation of the shirking dilemma. That dilemma captures some of the main features of such dilemmas, but it is not obvious that the results would extend to cases where the time spent shirking was essentially continuous (minutes or seconds) rather than discrete (hours or days) and when individuals are paid a wage rate equal to their marginal revenue product.

Solving shirking problems is important for the theory of firms, as noted in the discrete example. We'll again focus on a two-person team and use payoff functions from the multiplicative exponential functions used in Part I.

In this case, the payoffs are again linked to various combinations of leisure (for the shirker) and his or her share of the output of the team. We'll assume that the team exhibits modest increasing returns during the normal work day.

Suppose that Ellen (E) and Fred (F) work in a small business and engage in team production where the production function is of the form $Q = aW_E^b W_F^c$. (Note that the superscripts are exponents and the subscripts identify Ellen and Fred's hours of work.) The output is sold in a market at price P, and each team member is paid according to their marginal revenue product. We'll assume that the production function exhibits modest increasing returns over the course of the day, $1.5 > b > 1$ and $1.5 > c > 1$. (Otherwise, their wages would rise as they shirked.)

The neoclassical model of wage rates assumes that individuals are paid their marginal revenue product for each hour worked—but in this case there is a difference between hours on the job and hours worked.

They each spend 8 hours on the job, but they are paid for a full day (8 hours). Ellen's wage rate is $w_E = PQ_{W_E} = P\,(abW_E^{b-1}W_F^c)$ and Fred's income is similarly $w_F = PQ_{W_F} = P\,(acW_E^b W_F^{c-1})$. Each of the workers is supposed to work 8 hours, thus Ellen's income for the day is $Y_E = 8w_E = 8P(abW_E^{b-1}W_F^c)$, and Fred's income for such a day is $Y_F = 8w_F = 8P(acW_E^b W_F^{c-1})$.

We'll assume that neither is carefully monitored for their work-effort, and thus each may take part of the workday as leisure (shirk) without being punished by the firm's management.

Notice that each employee's marginal revenue product varies as the number of hours that their teammate works, as normally is the case for team production. Thus, their wage rates and income also vary. Each values both leisure and income from their jobs. Notice that insofar as they are paid their marginal revenue product, neither will want to work zero hours, because their marginal products and incomes would fall to zero in that case.

To simplify a bit, assume both Ellen and Fred have similar preferences for work and leisure. The utility realized from a day at "work" is $U_i = L_i^j Y_i^k$. where $L_i + W_i = 8$. Notice that the eight-hour day allows both the amount of leisure and income to be expressed as a function of the time spent actually working while on the job. After these substitutions, Ellen's utility function can be written as:

$$U_E = (8 - W_E)^j [8P(abW_E^{b-1}W_F^c)]^k \tag{13.1}$$

And for Fred, this is:

$$U_F = (8 - W_F)^j [8P(acW_E^b W_F^{c-1})]^k \tag{13.2}$$

To find out Ellen's ideal workday—time on the job actually spent working—differentiate equation 13.1 with respect to W_E, set the result equal to zero, and then solve for W_E.

$$-j(8 - W_E)^{j-1}[8P(abW_E^{b-1}W_F^c)]^k$$

$$+k(8 - W_E)^j [8P(abW_E^{b-1}W_F^c)]^{k-1}[8P(ab(b-1)W_E^{b-2}W_F^c)] = 0 \text{ at } W_E^*$$

or

$$j(8 - W_E)^{j-1}[8P(abW_E^{b-1}W_F^c)]^k =$$

$$k(8 - W_E)^j [8P(abW_E^{b-1}W_F^c)]^{k-1}[8P(ab(b-1)W_E^{b-2}W_F^c)]$$

Which simplifies to:

$$j[8P(abW_E^{b-1}W_F^c)] = k[8 - W_E][8P(ab(b-1)W_E^{b-2}W_F^c)]$$

And further to:

$$jW_E = k(8 - W_E)(b-1) = 8k(b-1) - k(b-1)W_E$$

or $$((j + k(b-1))W_E = 8k(b-1)$$

Which implies that

$$W_E^* = 8\left[\frac{k}{(\frac{j}{b-1}) + k}\right] \tag{13.3}$$

This is Ellen's best-reply function. She has a pure dominant strategy, because her best workday is not affected by Fred's choice. If b=1.25 and the other exponents were ½, this equation would imply that Ellen only works 1/5 of the time that she is "on the job."

As modelled, Ellen's shirking does not depend on Fred's effort, because of the functional form of the wage and utility functions assumed. This would not be true of all payoff functions, but is convenient here because a clear solution for the extent of shirking can be worked out. Fred would work the same number of hours if c=1.25:

$$W_F^* = 8\left[\frac{k}{(\frac{j}{c-1})+k}\right] \tag{13.4}$$

Neither pays attention to the fact that as they work less, the wage rate of the other falls and the team's output falls.

Of course, firms cannot long survive if they pay for work that is not actually being done as in this case, unless the team produces an extraordinary amount of output in a short time. It is to avoid such problems that firms monitor the productivity of their workforce and incentivize team members to work longer hours than most would have worked without the monitoring and incentives schemes developed.

This aspect of the productivity of firms is not modelled in the core neoclassical models, although it is one of their core functions. Note that salaries equal to marginal revenue products are not sufficient to solve the shirking problem.

Lottery-Like Contests with Stochastic Payoffs

The second contest with continuous strategy functions examined in this section is, like the prisoner's dilemma, a type of contest which is not itself of great relevance for economics—except that it has properties that are very similar to other choice settings of great interest.

It, like the multiplicative exponential functions used throughout the text, also has mathematical properties that generate clear results that are intuitively plausible and useful to demonstrate how rivalry in many stochastic winner-take-all contests and reward-sharing contests tends to operate.

The model developed characterizes the purchase of lottery tickets, where the probability of winning some price (or the share received) tends to increase as one purchases more tickets (e.g. engages in greater efforts) but falls as one's rivals purchase more tickets (e.g. engages in greater efforts). Lotteries, per se, are not central to the interests of microeconomists, but many choice settings that have these two properties exist and are of interest—as with competition in areas of research and development, marketing, and efforts to influence governments for favorable regulatory treatment

Any contest in which the prize is more or less fixed, and the odds of winning the prize increases with efforts have similar properties. In such settings, winning is not simply a matter of who buys the most tickets, as in an auction. Winning is a stochastic event. Even a person that purchases only a single ticket may win. Buying more tickets simply increases the probability of winning (as investing in R&D tend to increase the probability of a breakthrough patent, rather than assure it).

Consider, for example, a two-person lottery contest in which each can buy as many tickets as they want, and each player's probability of winning depends on the number of tickets owned by that player relative to the total sold. Both players are assumed to be risk neutral and thus maximize their "expected" net earnings from purchasing tickets.

Recall that the **expected value** of an event with outcomes 1, 2, 3, ... N is $V^e = \sum P_i V_i$, where P_i is the probability of event i, and V_i is the value of event i. In the contest of interest here, there are just two outcomes, winning and losing.

If Al purchases Na lottery tickets and Bob purchase Nb tickets, Al's expected profit (R_a^e) from purchasing lottery tickets is

$$R_a^e = [N_a / (N_a + N_b)]Y - N_a C \qquad (13.5)$$

where Y is the prize won, and C is the cost of a lottery ticket. Similarly, Bob's expected net benefit (profit) is $R_b^e = [N_b / (N_a + N_b)]Y - N_b C$.

Al's expected profit maximizing number of lottery tickets can be found by differentiating equation 13.5 with respect to Na and setting the result equal to zero.

$$\frac{dR_a^e}{dN_a} = \{[1 / (N_a + N_b)] - [N_a / (N_a + N_b)^2]\}Y - C = 0 \text{ at } N_a^* \qquad (13.6)$$

This expression can be solved for N_a^*. Putting terms over the same denominator and adding C to each side yields:

$$[N_a + N_b - N_a]/(N_a + N_b)^2 = C/Y \quad \text{or} \quad N_b/(N_a + N_b)^2 = C/Y$$

Solving this for N_a^* yields:

$$N_a^* = -N_b + (N_b Y/C)^{0.5} \qquad (13.7)$$

Equation 13.7 is Al's **best-reply function**. In this choice setting, it tells Al the expected profit maximizing number of lottery tickets to purchase given any particular purchase by Bob.

Since N_a^* varies with Bob's purchase, N_b, Al does **not** have a pure dominant strategy in this setting.

A similar best-reply function can be derived for Bob,

$$N_b^* = -N_a + (N_a Y/C)^{0.5} \tag{13.8}$$

The Nash equilibrium occurs when both persons are simultaneously on their best-reply functions. If both persons are simultaneously on their best-reply functions, then neither can change their strategy and improve their payoff (remember that the best-reply function for player i maximizes his or her payoff, given the strategies adopted by all other players), as required for the existence of a Nash equilibrium.

Thus, the Nash equilibrium of this lottery game occurs at a point where:

$$N_a^* = -N_b^* + (N_b^* Y/C)^{0.5} \text{ and } N_b^* = -N_a^* + (N_a^* Y/C)^{0.5}$$

To find the N_a^* and N_b^* combination where both these conditions hold, one can either substitute the equation describing N_b^* in terms of N_a^* into the Al's best-reply function and do quite a bit of algebra.

In a **symmetric game** (a game in which players have the same strategy sets and payoff functions) there is normally a symmetric equilibrium. In that case, the two best-reply functions will intersect at a point where $N_a^* = N_b^*$.

Using this principle allows us to greatly reduce the algebra necessary to solve for the Nash equilibrium purchases of lottery tickets. At the symmetric lottery game's equilibrium:

$$N_a^* = -N_a + (N_a^* Y/C)^{0.5}$$

$$\text{or } 2N_a^* = (aY/C)^{0.5}$$

Squaring both sides, we have: $4N_a^{*2} = N_a^* Y/C$ which implies that $4N_a^* = Y/C$

$$N_a^{**} = Y/4C \tag{13.9}$$

and since $N_a^{**} = N_a^{**}$ at the symmetric Nash equilibrium, we also have $N_b^{**} = Y/4C$.

Each ticket costs C, thus Al spends $Y/4$ on tickets—1/4 of the prize on offer. The same is true for Bob, so it is clear that this particular lottery will not be a "money maker" for its organizers. Together Bob and Al spend only half the amount promised as a prize. Nonetheless, a good deal is invested relative to the value of the ultimate prize.

Within political economy, a topic taken up in part III of the book, the most common use of the lottery model is to characterize political rent-seeking games, an application originally developed by Gordon Tullock (1980).

Generalizing the Lottery Contest to N Players

The lottery game and its various applications can easily be generalized to take account of more than 2 players, and to include "technologies" where the exponents on investments are subject to increasing or decreasing returns.

Suppose that there are N players. Let K represent the number of tickets purchased by all the players other than player A. In that case, the expected payoff of a "typical" player (here player a) can be written as:

$$R_a^e = [N_a / (N_a + K)]Y - N_a C \qquad (13.10)$$

Differentiating with respect to Na yields:

$$\frac{dR_a^e}{dN_a} = \{[1 / (N_a + K)] - [N_a / (N_a + K)^2]\}Y - C = 0 \; at \; N_a^*.$$

Solving for Na, as above, yields:

$$N_a^* = -K + (KY/C)^{0.5} \qquad (13.11)$$

This equation is the best-reply function of a typical player in the present N person game, which characterizes his or her best response to the total efforts (number of tickets) of the other players.

At the symmetric equilibrium, all players purchase the same number of tickets, thus,

$$K = (N-1)N_a^{**}$$

This allows equation 13.11 to be rewritten as:

$$N_a^* = -(N-1)N_a^{**} + ((N-1)N_a^{**}Y/C)^{0.5}$$

Solving for Na*, yields:

$$N_a^{**} = [(N-1)/\mathrm{N}^2]\,(Y/C) \qquad (13.12)$$

Note that when N = 2, as above, $N_a^{**} = \left[\frac{2-1}{(2)^2}\right]\left(\frac{Y}{C}\right) = \left(\frac{1}{4}\right)Y/C$, as before.

The **total expenditure** on lottery tickets is NC times this amount, or $[(N-1)/N]Y$. Note that total expenditures approaches Y in the limit as N approaches infinity. In this case, a lottery is a breakeven proposition for the organizers (assuming that the ticket buyers are risk neutral and have no charitable inclinations).

Note also that if this contest represents a R&D contest for a patent that is worth Y, the researchers collectively spend as much as the prize, although the firm that makes the breakthrough and gets the patent, has earned a very substantial return on its investment.

Different technologies for increasing one's chance of winning can also be taken into account by assuming changing our assumptions about investments in the game (the number of tickets bought, Na) affect the probability of winning the prize. For example, we can take account of economies and diseconomies of scale by changing from

$$P = \frac{N_a}{N_a + K},\ to\ P = N_a^d / (\sum N_i^d).$$

The payoff function for a typical player now becomes:

$$R_a^e = [N_a^d / (\sum N_i^d)]Y - N_a C \qquad (13.13)$$

Differentiating with respect to N_a now yields:

$$[\frac{dN_a^{d-1}}{(\sum N_i^d)} - \frac{(N_a^d)(dN_a^{d-1})}{(\sum N_i^d)^2}]Y - C = 0$$

At the symmetric equilibrium, $N_a^{**} = N_i^{**}$ for all i = 1, 2, ... N, thus:

$$[\frac{dN_a^{d-1}}{(NN_a^d)} - \frac{(N_a^d)(dN_a^{d-1})}{(NN_a^d)^2}]Y - C = 0$$

Putting the numerators over a common denominator and collecting a few terms yields:

$$[\frac{NN_a^d dN_a^{d-1}}{(NN_a^d)^2} - \frac{(N_a^d)(dN_a^{d-1})}{(NN_a^d)^2}]Y - C = 0$$

, or

$$\left[\frac{d(N-1)N_a^{2d-1}}{N^2 N_a^{2d}}\right]Y = C$$

Solving for Na*, yields the individual's number of tickets (level of resources invested in the contest) at the symmetric Nash equilibrium:

$$[d(N-1)N_a^{2d-1}]Y = CN^2N_a^{2d}$$

$$N_a^{**} = [(N-1)/N^2]\,(dY/C) \quad (13.14)$$

Note that when $d = 1$ and $N = 2$, as above, $N_a^{**} = (1/4)\,(Y/C)$, as before.

The **total expenditure** on "rent seeking" is again NC times this amount.

$$N(N_a^{**})C = d[(N-1)/N]Y \quad (13.15)$$

Note that total expenditures can now exceed Y, if $d(N-1)/N > 1$.

The implications of our analysis of contests resembling lottery games can be summarized as follows:

i. The more players are in the game, the less each player spends.

ii. Nonetheless, the total spent rises with the number of players.

iii. In games with constant returns (the classic Tullock contest function) the total investment in the contest approaches the value of the prize (Y) as the number of players approaches infinity.

iv. Contests with increasing returns may have "super dissipation," where more resources are invested in the contest than the prize to be won.

v. (Note that no player will routinely play such games. However, "no one" playing is also not an equilibrium, so potential players may play mixed participation strategies. Mixed strategies are discussed in the appendix.)

As mentioned above, there are a surprisingly large number of applications of these lottery-types of games. Essentially any contest in which one's own additional resources increases the probability of winning while that of one's rivals reduces the probability of winning can be modeled with such functions. Similarly, any contest in which additional resources increases one's share of a prize, while expenditures by one's rivals reduce one's share, can be modeled in the same way.

V. Characterizing Contests Using Very Abstract Payoff Functions

More general characterizations of games are also possible using abstract functions to characterize payoff functions and strategies. For example, the payoff function of a two-person game analogous to the lottery game can be written as can be represented using abstract functions.

For example, let the payoff of player A be $G^A = g(X^A, X^B)$ and that of player B be $G^B = g(X^B, X^A)$ where X^A is the strategy chosen by player A and X^B is the strategy chosen by player B.

Each player in a non-cooperative game attempts to maximize his payoff, given the strategy chosen by the other. To characterize the payoff maximizing strategy for player A, assume that the payoff function is twice differentiable and differentiate A's payoff function with respect to X^A and set the result equal to zero. X^{A*} has the property that

$$\frac{dG^A}{dX^A} = 0. \qquad (13.16)$$

The implicit function theorem implies that A's best strategy X^{A*} is a function of the strategy chosen by player B, $X^{A*} = x^A(X^B)$. This is A's best reply function. A similar reaction (or best-reply) function can be found for player B.

At the Nash equilibrium, both reaction curves intersect, so that:

$$X^{A*} = x^A(X^{B*}) \text{ and } X^{B*} = x^B(X^{A*}) \qquad (13.17)$$

In a symmetric game, the Nash equilibrium simply requires: $X^{A*} = x^A(X^{A*})$. Which implies that the Nash equilibrium is the fixed point of each player's identical reaction function.

Comparative statics can be undertaken with the implicit function differentiation rules in a manner similar to that applied with respect to the market equilibria modeled in Chapter 5 as long as the payoff function is strictly concave and twice differentiable.

VI. Conclusion

Applications of game theory in economics tend to focus on relatively small number settings where an economically relevant outcome is jointly determined by the individuals, groups, or organizations making the decisions that ultimately determine the outcome of interest. Such settings include various duopoly and oligopoly models of markets, relationships between firms and their employees and shareholders, innovation contests, and marketing contests—to name just a few of the applications that economists have found of interest during the past half century. Within political economy, contests among governments, among interest groups that attempt to influence policy decisions, and among politicians and political parties for votes in elections have also attracted considerable attention.

In some cases, game theory allows previously unnoticed problems to be analyzed and thus provides explanations for some of the standing procedures that are otherwise difficult to account for. In some cases, they also have clear implications about how existing procedures can be improved. In others, new light is shed on older questions such as the extent to which firms may compete with each other in various market types and with each other for various government favors and preferences.

The overall effect is to enhance our understanding of a large number of small-number relationships within markets and firms that are largely ignored by the core neoclassical models of market prices.

Appendix: Some Technical Game Theoretic Terms and An Existence Proof

I. Useful Game Theoretic Terminology

Definition: A *normal or strategic* form game is given by: (i) a list of players i = 1, ...I; (ii) a list of strategies Si that player i might employ; (iii) a payoff function which defines the payoffs realized by each player under all possible combinations of strategies. (See Kreps p. 379)

Strategies may be complex in the sense that they involve a sequence of moves or conditional play, or even involve random play.

Games in normal form are often represented with matrices (as we did above for the Prisoner's Dilemma game) or represented mathematically with payoff functions for each player in each possible strategic circumstance. (As we did above for the Cournot Duopoly problem) A sequence of play can be represented as a vector of moves. Every game in extensive form can be represented in normal form.

Definition: A game in *extended form* is represented as a *game tree* with (i) a list of players, (ii) an assignment of decision nodes to players or to nature, (iii) lists of

> actions available at each decision node, (iv) and a correspondence between immediate successors of each decision node and available actions, (v) information sets, (vi) an assignment of PAYOFFs for each player at each terminal node (possible ending of the game) (vii) and probability assessments over the initial nodes and over the actions at any node that is assigned to nature. (Kreps, p. 363)

A game in extensive form represents game play as a sequence of choices made by game players. "Simultaneous" choices are represented with restrictions in the information set. If one player does not know what the other chose at the previous "node(s)," then it is as if the whole sequence of decisions is made simultaneously by all players. Games in extensive form are convenient ways to represent games where a sequence of moves are made (as in chess or checkers) as well as games where information changes in systematic ways as the game unfolds. *Any game in Normal form can be represented as a game in extensive form.*

> **Definition**: A *mixed strategy* is a random strategy. Suppose that the range of possible actions (or sequences of actions) is S. Then a probability function defined over S is a mixed strategy.

Note that a mixed strategy does not have to assign non-zero probabilities to all of the possible actions that may be taken. One can consider a *pure strategy* to be a special case of a mixed strategy where all the probabilities of actions (or sequences of actions) are zero except for one (the pure strategy) which is played with probability of 1.

In games with no equilibrium in pure strategies, as in the "even or odd finger game," or "paper, scissors, rock" game, most people intuitively use something like a mixed strategy. That is, they vary their actions without obvious pattern. To be a mixed strategy, the pattern must be determined by a particular probability function (chosen by the individual players).

II. Sufficient Conditions for the Existence of Nash Equilibria in Non-cooperative Games

Proposition. *Every finite player, finite strategy game has at least one Nash equilibrium* if we admit mixed strategy equilibria as well as pure. (Kreps p.409 and/or Binmore p.320).

The proof relies upon Kakutani's fixed point theorem, which is a generalization of Brouwer's fixed point theorem used in the General Equilibrium existence proof derived above. Here is a condensed version. (Presented in Kreps, page 409).

Let i = 1, ... I be the index of players, let Si be the (pure) strategy space for player i and let Σi be the space of probabilities distributions on Si.

The strategy space of mixed strategy profiles is $\Sigma = \prod_i^I S_i \, \Sigma_i$, which is the cross product of all individual mixed strategies.

For each combination of mixed strategies $\sigma = (\sigma_1, \sigma_2, \ldots\ldots \sigma_i)$ find each person's best-reply function for player i given the other strategies, $\varphi_i(\sigma^{\sim i})$.

Define $\varphi = (\varphi_1\ \varphi_2 \ldots\ldots \varphi_i)$ which is vector of best-reply functions. Note that φ is a mapping from the domain of mixed strategies onto itself. It is upper semicontinuous and convex, hence by Kakutani's fixed point theorem a *fixed point exists*.

This fixed point is a mixed strategy I-tuple which simultaneously characterizes and satisfies each player's best-reply function. That is, it is a Nash equilibrium. Q. E. D.

A somewhat less general, but more intuitive proof is provided by Binmore in section 7.7.

In relatively simple games, it is often possible to compute the Nash equilibrium in mixed strategies by adjusting player 1's probabilities to make player 2 indifferent between his or her strategies, and player 2's probabilities to make player 1 indifferent between his or her available strategies.

In complex games, like chess with a time constraint, one can prove that an ideal mixed strategy exists but cannot calculate it. Fortunately, a good many economic games have computable equilibria with pure strategies (duopoly, monopolistic competition).

III. Subgame Perfection

Definition: A **subgame perfect Nash Equilibrium** for game in extensive form is a Nash Equilibrium of the game that is also a Nash equilibrium in every proper subgame of the game.

Definition: A proper subgame of an extensive game is a node t and all its successors.

That is to say, the choice called for at node "t" at the beginning of the game is the same choice that would have been chosen had the game started at node t-1. Sub-perfect equilibrium, thus implies that a strategy is "self-enforcing" in the sense that if such a strategy exists for each player, they will play out the whole series of moves called for in the strategy chosen in the first period (their equilibrium strategic plan).

Subgame perfection is an important property of *self-enforcing contracts and credible commitments*. Under a self-enforcing contract, each participant has an incentive to abide by all of their contractual obligations, even if there are multiple opportunities to renege. This equilibrium concept has been applied to insurance problems, to labor markets, political constitutions, and international law (where there is no supra-national law enforcing agency).

Another method is to assure that each round of the game involves dominant strategies (with mutual gains to trade). In this case there is no "hold out" or end-game problem, because it is in each person's interest to perform their "duty" in each period.

Illustrations of Sub-Game Perfection

Proposition. Mutual cooperation in a repeated PD game with a known end point is never subgame perfect.

Proof using backward induction.

In the last round one is in an unrepeated PD game. Consequently, the Nash equilibrium is for mutual defection.

If there is no cooperation in the last round, then one cannot lose future cooperative advantage from failing to cooperate in the second to last round.

Similarly, if both players will not cooperate in the second to last round, there is no risk (of retaliation) to defecting in the third to last round and so on.

"Rational" players will never cooperate in a repeated PD game of finite length with a known end point. [This may be good news for markets but not for collective action.]

Note that this may not be true of finite games with an unknown end point. Why?

A similar problem is associated with what is called the ***Centipede Game***: a finite length game in which each player may choose to stop the game on his turn, in which case he or she "wins" the entire score, however he also knows that scores will increase if play continues beyond the other player(s) turn. At time T, the game ends.

Clearly, if any player can do better stopping before time T, he will stop at that earlier time, but that implies a "new" shorter game with a similar problem for all. This causes some other player to end the game still earlier and so on. **Stopping on the first turn** winds up being the only subgame perfect equilibrium in the standard form of the Centipede Game.

IV. The Folk Theorem

The *folk theorem* applies to repeated games, especially games that continue forever or have an unknown (random) end point. Essentially the folk theorem says that essentially all possible patterns of strategies can be equilibria to a game, if players are able to make credible threats.

For example, one of many possible equilibria is mutual cooperation in a repeated PD game. The intuition is the following: Suppose that in a repeated PD game both players announce that if the other one testifies, he will testify against the other in every successive PD game. This announced strategy, *if believed*, eliminates the "temptation" payoff in the off diagonal cell. The cumulative payoff from repeated play now exceeds the one time off-diagonal payoff followed by a long (possibly infinite) series of non-cooperative payoffs. So, under these announced strategies, mutual cooperation is the best strategy.

With a suitable adjustment in probabilities, various equilibria in conditional mixed strategies can be found which result in average payoffs a bit higher than the mutual defection payoff up to the complete cooperation payoff.

The troubling thing about the folk theorem, is that it allows too many outcomes to be equilibria, so it has little predictive value. It demonstrates that the equilibrium of a sequential game depends entirely on the announced or predicted **conditional** strategies of the players in the game, and the **credibility** of those announcements or predictions.

It is clear that a strategy that is subgame perfect is a more credible threat than one that requires "non-rational behavior." Consequently, equilibria strategy pairs under the folk theorem should be subgame perfect, if they are to be plausible combinations of strategies.

One method of inducing particular strategies is to post mutual bonds (or hostages) with third parties which are valued by each hostage by more than the value of "defecting" rather than "cooperating." In such cases, a neutral third party (non-player) returns the hostage 1 to party 1 only if party 1 has executed his contractual obligations. The hostage/bond method induces cooperation by changing the payoffs of the original game, and it is credible as long as the bonding agent benefits from returning the hostages.

V. A Few Review Problems

Stackelberg Game. Suppose that Acme and Apex are duopolists who have identical total cost functions, C = 100 + 10Q, and "share" the same market (inverse) demand curve: P = 1000 - 20Q.

i. Acme makes its output decision first. What output should it choose if it knows that Apex will simply maximize its own profits given Acme's output? ii. Is the resulting market equilibrium sub-game perfect? Explain.

iii. Does the Stackelberg equilibrium differ from the Cournot equilibrium for this pair of firms? Demonstrate and explain.

Continuous Dealings. Suppose that Al uses George's garage for all car repairs. Al tells George that if he ever believes he has been cheated by George, he will never return.

Suppose further that, ex post, cheating can always be determined by Al. George gains $25.00 each time he honestly services Al's car and $50.00 if he cheats. If Al leaves before service is obtained his payoff is 0. Al receives $15.00 of consumer surplus if he uses George and $5.00 if he uses another garage (known to be honest, but a bit further away and more expensive). However, if George cheats, Al loses $15.00 (of surplus).

Analyze this as a one-shot game. Should George cheat Al and/or should Al use George? Explain.

Now consider the setting in which the garage game is to be repeated *ad infinitum*. Is the game now subgame perfect in non-cheating by George and use of George's by Al? Demonstrate and explain.

VI. Selected References

Binmore, K. G. (2007). *Game Theory: A Very Short Introduction*. Oxford UK: Oxford University Press.

Dixit, A. K., Skeath, S., & McAdams, D. (2025). *Games of strategy* (6th ed.). W. W. Norton & Company.

Kreps, D. M. (1990). *Game Theory and Economic Modelling*. Oxford UK: Oxford University Press.

Luce, R. D. and Raiffa, H. (1957). *Games and Decisions, Introduction and Critical Survey*. New York: Wiley.

Murdoch, J. C., & Sandler, T. (1997). Voluntary cutbacks and pretreaty behavior: The Helsinki Protocol and sulfur emissions. *Public Finance Review, 25*(2), 139–162.

Owen, G. (2013). *Game Theory*. Bingley UK: Emerald Group

Skyrms, B. (2004). *The Stag Hunt and the Evolution of Social Structure*. Cambridge UK: Cambridge University Press.

Tullock, G. (1980). Efficient rent seeking. In J. M. Buchanan, R. D. Tollison, & G. Tullock (Eds.), *Toward a theory of the rent-seeking society* (pp. 97–112). College Station: Texas A&M University Press.

Vanberg, V. J., & Congleton, R. D. (1992). Rationality, morality, and exit. *American Political Science Review, 86*(2), 418–431.

Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton NJ: Princeton University Press.

Weintraub, E. R. (Ed.). (1992). *Toward a History of Game Theory* (Vol. 24). Duke University Press.