An Introduction to Non-Cooperative Game Theory

The use of game theory to analyze economic problems is the most active area of theory in modern economics. A quick look at any economics journal published (and many other social science and philosophy journals) in the past two decades will reveal a large number of articles that use elementary game theory to analyze economic behavior in a variety of institutional and/or natural settings. Modern work on self-enforcing contracts, credible commitments, the private production of public goods, externalities, time inconsistency problems, models of negotiation, and models of political and social activity all use game theoretic models as their "engines of analysis."

Although formal game theory is relatively new, it can be argued that the use of game theory in economics is older than game theory itself. For example, the Cournot duopoly model (1838) is an early example of a non-cooperative game with a Nash equilibrium. Stackelberg's (1934) analysis of duopoly is also game theoretic as are essentially all models of monopolistic competition. Game theory as a field in applied mathematics emerged after World War II.

Most economic applications of game theory apply the rational choice models (utility maximizing or net-benefit maximizing models) that we have used throughout this course. That is to say, the game theorists normally assume that players are "rational" in that the selection of strategies can be modeled as utility maximizing or net benefit maximizing choices. The result of the strategy choices of all the players in the game determine the payoffs (utilities or net benefits) actually received by the players.

This chapter provides a short introduction to game theory. We'll use rational choice models to analyze economically relevant behavior in various choice settings. Most of those settings analyzed involve relatively simple choice settings in which there are just two players and two possible strategies. However, as true of many simply models in economics, much can be learned by understanding how the payoffs of choice setting (incentives) affect behavior and consequences of that behavior. Extensions of the simpler models involve shifts to settings with more possible strategies, larger numbers of players, and repeated games. Some of these are undertaken in the next chapter.

I. Non-Cooperative Games.

A. In "non-cooperative games" individuals are normally assumed to maximize their own utility without caring about the effects of their choices on other persons in the game. There is normally little or no communication among players and the games that attract the most attention (although they are not necessarily the most important) are those the generate outcomes that are Pareto suboptimal.

- In such cases, the outcome of the game is such that a change in strategies could have made at least one person (player) better off without making anyone worse off.
- This class of choice settings are sometimes termed social dilemmas.

i. The outcomes of most non-cooperative games are jointly determined by the strategies chosen by all players in the game.

- Consequently, each person's welfare depends, in part, on the decisions of other individuals "in the game."

ii. Normally the payoffs are in terms of "utility," "euros," or "dollars," but occasionally other values are used.

iii. The assumption that players do not know what the other is going to choose allows simple two person games to be characterize choices in other settings with large numbers of players, because it is unlikely that a large number of players will know each other very well or be able to communicate with all the others.

ii. Examples of economically relevant games include:

- Voluntary exchange can be characterized as a game in which one player makes offers to sell and the other makes offers to purchase. The case in which all gains from trade are realized is one of the outcomes and is not always, as it turns out, the most likely to emerge.
- Cournot duopoly, a choice setting where the profits of two (or more) firm's depend upon their own output decisions and the output decisions of the other firm(s) in the market.
- Public Goods Problems: After a snow fall, the amount of snow on neighborhood sidewalks depends partly on one's own efforts at shoveling and partly that of all others in the neighborhood.
- Democratic Politics: In an election, each candidate's vote maximizing policy position depends in part on the positions of the other candidate(s). The policies chosen by the candidate that wins the election often have economic consequences.

iv. Game theory models are less interesting (and usually less useful) in cases where there are no interdependencies.

- For example, a case where there is no interdependence it that of a producer or consumer in perfectly competitive market.
- In that setting, a consumer (or firm) is able to buy (or sell) as much as they wish without affecting market prices.
- Game theory can still be used in such cases, but with little if any advantage over conventional tools.

B. The choice settings in which economists most frequently apply game theory are small number of player settings in which outcomes are jointly determined by players decisions that are made independently of one another.

C. A game is said to be completely specified, when the players participating it and the payoffs of every possible combination of strategies is clearly defined (or characterized).

D. The simplest game that allows one to model social interdependence is a "symmetric" two-person "one shot" game in which both players independently choose between two strategies, S_1 and S_2.

- In a semmetric game, players can use the same strategies and have essentially the same payoffs: identical payoffs when both use the same strategy and mirror image payoffs in cases in which they do not.

i. There are four possible outcomes to such games:

- (1) both players may choose S_1,
- (2) both may choose S_2
- (3) player A may choose S1 and player B may choose S2,
- or (4) vice versa.

ii. **At a Nash equilibrium, no player can independently change his or her strategy in a manner that increases his or her payoff.**

iii. The particular combination of strategies that emerges at the equilibrium is the result of the independent optimizing decisions of the two players, A and B (Al and Bob).

E. For example, consider the "trading game" with transactions costs to the right below

i. Bob has cash and Al has apples. Bob is thinking trading some cash for some of Al's apples. Al is thinking about trading apples for some of Bob's cash.

ii. Since trade is voluntary, nothing happens unless both players agree to trade.

iii. However, for the purpose of illustration, It is assumed that it costs "one util" a bit to make an offer, whether taken or not.

iv. Thus, the lower left hand and upper right-hand cells have payoffs for Al and Bob that are lower for the one making the offer (trading), while the other is unaffected.

Bob

	Seek Trade	Don't
AL Offer Trade	(a, b) (10,10)	(a, b) (-1, 0)
Don't	(0,-1)	(0,0)

v. The off-diagonal cells reflect the transactions cost of the attempting the exchange on either the seller or buyer side of the market.

vi. Note that both the upper left and lower right can be Nash equilibrium.

vii. (The upper left is the textbook equilibrium, the lower right is the no-trade equilibrium, where you leave potentially tradable or saleable products in your closet rather than put them on Craig's list or shop on Craig's list.)

F. A game can be said to have a **Nash Equilibrium** whenever a strategy combination is "stable" in because no player in the game has an interest in changing his or her strategy, given that of other players. Neither player can increase his or her payoff by doing so—given what the other has chosen.

i. Note that 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) given that of the other player(s) in the game in these cells (strategy combinations).

ii. A state of the world or game outcome is said to be **Pareto Optimal** or Pareto Efficient, if it is impossible to reach another state (outcome) where at least one person is better off, and no one is worse off.

- Note that the (Trade, Trade), equilibrium is Pareto optimal, but none of the other possible outcomes are.
- **DEF**: An outcome is Pareto Efficient if and only if there are not Pareto superior moves from that outcome.
- **DEF:** State X is Pareto superior to state Y, if and only if at least one person prefers X to Y and no one prefers Y to X.

iii. What may be surprising is that 2 person - 2 strategy games can often capture the essential features of particular choice settings of interest to social scientists.

- This is also true of the Edgeworth box, which very nicely illustrates the gains from trade that potentially can be realized by both sellers and buyers, even in barter economies.

- However, as the above game theoretic representation of the "problem of exchange" demonstrates, trade is not as automatic as the Edgeworth box and textbook economics implies.)

II. The Prisoners' Dilemma Game: Another Simple Non-Cooperative game.

A. The *Prisoners' Dilemma game* is probably the most widely used game in social science.

i. PD games are normally "symmetric games" because the players have the same strategy sets and payoff functions.

ii. There are both general PD games and specific ones such as the original prisoner's dilemma game

iii. The "original" prisoner's dilemma game goes something like this: Two individuals are arrested under suspicion of a serious crime (armed robbery). Each is known to be guilty of a minor crime (say shoplifting), but it is not possible to convict either suspect of the serious crime (armed robbery) unless one or both of them testifies that the other is guilty of the more serious crime. The police separate the prisoners and tell both that they will receive a lower penalty if they testify against the other

iv. The payoffs listed in the game matrix are **years in jail** or monetary penalties rather than utility in this case.

Prisoner A \ Prisoner B	Testify	Don't
Testify	(<u>10</u>,<u>10</u>)	(<u>1</u>, 12)
Don't	(12,<u>1</u>)	(2,2)

v. The prisoners are separated. Each is told that if he testifies about the other's guilt that he will receive a reduced sentence for the crime that he is known to be guilty of.

vi. The Nash equilibrium of this game is that BOTH TESTIFY against the other.

B. **To see why this equilibrium emerges**, analyze the behavior that is likely to emerge given the payoffs represented in the matrix.

i. Each cell of the original Prisoner's dilemma game matrix contains payoffs for A and B, in **years in jail** (a bad).

- **Each individual prisoner will rationally attempt to minimize his jail sentence.**
- Note that regardless of what Prisoner B does, Prisoner A is better off testifying. 10 < 12 and 1 <2. **Testifying is the dominant strategy**. (Notice that I've underlined Al's best response to the strategies that might be chosen by Bob, and also done so for Bob's best response to the strategy choices of the Al.)
- Note that the same strategy appears to yield lower sentence for Prisoner B. If A testifies, then by also testifying B can reduce his sentence from 12 to 10 years. If A does not testify, then B can reduce his sentence from 2 to 1 year by testifying.

ii. When each player in a PD or other games has a **dominant strategy**, the Nash equilibrium is easy to find—both players adopt the same strategy (their dominant strategies).

- The (testify, testify) strategy pair yields 10 years in jail for each.
- {Note that PD games have a single Nash equilibrium, but other games (such as the trading game above) may have more than one Nash equilibria. The uniqueness of the equilibrium is a consequence of their (pure) dominant strategies.)

iii. A more general method for identifying the Nash equilibrium associated with a particular game matrix is to underline that best strategy given particular choices by the other player. This process traces out a player's "best reply function." The same process can be undertaken for each player. **The equilibrium occurs where the two best reply functions intersect, which will be in cells in which both payoffs are under lined.**

iv. The Nash equilibrium of PD games is a "**dilemma**," because each prisoner would have been better off if neither had testified. (2 < 10).

- Thus, the PD game demonstrates that independent rational choices do not always achieve Pareto optimal results.
- Note that this equilibrium is good for society but bad for the criminals—assuming that both are actually guilty of the more severe crime. (But, this is not usually considered a property of PD games, in general.)

C. It turns out that the payoff structure of the prisoner's dilemma game (PD game) can be used to illustrate a wide range of social dilemmas.

i. Competition between Bertrand (price setting) duopolists.

ii. Competition between Cournot (quantity adjusting) duopolists.

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

iv. Competition among students for high grades in universities

v. Commons Problems

vi. Public goods problems

vii. The Not In My Backyard (NIMBY) Problem

viii. The free rider problem of collective action.

ix. The international regulation dilemma

x. The arms race

xi. The Hobbesian Jungle

xii. A variety of team production problems (shirking)

xiii. Contract Breach/Fraud (in settings with and without penalties)

- A few of such applications are illustrated below.

D. 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.

i. This contrasts with independent choices in markets, where independent decision making tends to make the participants better off.

ii. For the "defect, defect" outcome to be a unique equilibrium, defecting has to yield a payoff that is a bit higher than the cooperative solution regardless of whether the other player cooperates or not.

iii. PD payoffs are often represented "ordinal" utility levels with (3, 3) for the mutual cooperative solution and (2, 2) for the mutual defection result. In such cases, the other payoffs are often represented as (1,4) and (4,1) with the defector receiving 4 and the cooperator 1.

- In lectures and most other game matrices, I normally use utility payoffs since they tend to be more intuitive to econ students than penalties—because "more is better" rather than less is better. But, the "original" PD game was described in terms of years in jail and arguably is easier to interpret than the game in the utility levels associated with different amounts of jail time.)
- Tangible consequences, rather than utility can also be used (for example, years in jail for the original PD game, or outputs from the placing cattle on a commons for the commons problem).
- In most cases, high numbers are good rather than bad, unlike the original PD game.

iv. The PD payoffs can also be represented algebraically using (abstract) payoffs.

- (C, C) and (D, D) are the payoffs of the mutual cooperation and mutual defection outcomes
- Let S be the "sucker's payoff" and T be the "temptation payoff. (S, T) and (T, S) are thus the payoffs when one person defects and the other is "played for a sucker.
- In a generalized PD game, T>C>D>S.

E. 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).

- However, these assumptions can be dropped without changing the basic conclusion of the analysis.
- Essentially the same conclusions follow for N-person games in which the players have a infinite number of strategies and play for any finite number of rounds, as we will see later in the course.

F. Another limitation of PD games is that they assume that players have to play them. They cannot simply stop playing.

i. However, the possibility of exit can also be analyzed with game matrices.

ii. Prisoner dilemma games with exit (PDE games) were examined by Vanberg and Congleton (1992) for cases in which not all players were "rational" in the sense that they always play the defect strategy when they participate in the game. However, it turned out that "nicer" players were rational in another sense. If there are other players who also deviate from the defect strategy, cooperative enterprises will last longer than if there are not.

iii. Vanberg and Congleton (1992) showed that in PDE environments, strategies of conditional co-operation could be viable if the game is repeated many times.

- In PDE settings, relatively "nice guys" often first rather than last.
- Using norms favoring cooperation in PDE settings can be viable strategies for individuals, rather than one that requires personal sacrifice, as long as there are at least a few other "nice" players in the population of players that one might be matched with.
- Exit turns out to be a powerful constraint on exploitative players.

III. A Few Other "Named Games"

A. Note that the mathematical requirements for completely specifying a game are met in the Prisoner's Dilemma game.

i. The possible strategies are completely enumerated, and the players are clearly enumerated.

ii. The payoffs for each player are completely described for all possible combinations of strategies.

iii. The information set is (implicitly) characterized. (A player is said to have perfect information if he knows all details of the game. A perfectly informed player knows the payoffs for each party, the range of strategies possible, and whether the other players are fully informed or not.)

B. Several other interesting games are widely enough used to have been given names. Most can be created by changing the payoffs of the two-player two-strategy games.

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.

- (Individuals with no training in economics seem to regard all economic activities as zero sum games. Of course, in most cases, exchange creates value for each player. Trade is a positive sum game, not a zero sum 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 that a convention or norm is followed by both persons. Both players benefit if both do either strategy C or strategy D.
- For example, if we all drive on the left side of the road or all on the right, we all have higher payoffs than some drive on each side of the road.
- (Draw a sample game matrix for a coordination game and determine its Nash equilibria.)

iii. **Assurance games** are similar to coordination games. The off diagonal payoffs for the "cooperative" strategy are equal to or below those of the on diagonal cells (2, 0), however the upper left-hand "cooperative" cell has a higher value to both players (3, 3) than the lower right-hand "do nothing" cell, the original position (2, 2).

- It will take some kind of guarantee or trust to generate moves from the original lower right-hand score to the higher upper left-hand cell.
- {The trading game developed above is an assurance game.
- (Draw a sample game matrix for an assurance game and determine its Nash equilibria.)

iv. Some game theorists have renamed the **assurance game** a stag-hunt game after a setting described by Rousseau in his Discourse on Income 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' (2001) 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, of course, stable!)

v. C**hicken** games are contests in which identical strategy choices are disastrous rather than beneficial.

- In Chicken games, the "on-diagonal" strategies are lower than the “off-diagonal outcomes” yield higher payoffs.
- 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”.
- The off-diagonal payoff of the chicken is lower than that of the person who stays in his or her original lane, but normally higher than that received in the on-diagonal cells (e.g. the ones with both drivers do the same thing and crash).

- The off diagonal scores are generally higher, although one person does better than the other as with (4,2) and (2,4).
- (Create an illustrating payoff matrix for the dangerous 1950s teenage drivers game of chicken on rural roads.)

C. A Few Social Science Applications of PD-like Game Matrices

i. As noted above, a variety of social dilemma problems can be analyzing Prisoner's Dilemma Games.

ii. One such game is the "Public Goods" or free rider problem.

iii. In this game, a public service can be produced by either player alone by paying the full cost of the service, or it can be jointly produced if each pay's for half of the service.

- (A **pure public good** is a good that is "perfectly shareable," a good which once produced can be enjoyed by all in the community of interest.)

iv. (There are two versions of this public goods game and I will probably develop the other form in lecture. Their logic and outcomes are pretty similar, but the choice settings are a bit different. The one used below is often used in public goods experiments.)

v. Suppose that the benefits from production of the service takes is simply $(C_A + C_B)P$ where C_A is Al's contribution, C_B is Bob's contribution and P is the productivity of the contributions. The net benefit is simply the benefit from the public service less the player's contribution.

- If both players contribute C to the cost of the public good, then both players receive a **net benefits** (payoff) equal to 2CP - C.
- If neither player contributes, then the payoff for each is 0.
- If only one contributes, the person contributing receives PC-C and the person free riding receives PC.
- This structure of contributions and benefits yields a game with the following net benefit payoffs::

A Public Good Game

		Bob	
		Contribute	Free Ride
Al	Contribute	(A, B) (2CP-C, 2CP-C)	(A, B) (CP - C, CP)
	Free Ride	CP, CP - C	(0,0)

vi. Note that the magnitude of P (the productivity of contributions) is critical to whether there is a public good problem or not.

- For there to be a problem, CP>2CP-C>0.
- A bit of algebra allows one to show that CP>2CP - C > 0 requires 1>P>0.5.

- (Prove this as a practice exercise.)

vii. When P is in this range, the result is the classic Public Goods problem studied in Public Economics.

- {(Note that the payoffs in this case have the same rank order as those in a PD game, but in this case our simple production function motivates the payoffs. They are not simply assumed, but modeled.)
- {The Nash Equilibrium of a game where 1<P<0 is mutual free riding, which generates the (0,0) outcome.

viii. This example shows how a more realistic (detailed) choice setting can be represented using a game matrix.

- Of course, in most cases, more than two persons will be involved in paying for or producing the public good, and in most cases the good itself can be produced at various levels.
- Still, the 2x2 representation, captures essential features of the "free rider" problem that must be confronted when thinking about the production of public goods.

ix. It is interesting to note that if P> 1, no dilemma exists and the public good will be voluntarily provided.

D. The same matrix can be modified slightly to show how rewards and penalties can be used to solve such problems.

i. Assume that 1 < P < 0, in the above game.

ii. Suppose that free-riding can be observed, and that penalty F is imposed on any one that free rides.

- The penalty could literally be a fine or a tax imposed on free riders.
- The penalty might also be nonpecuniary, as with losing the respect of approval of one's friends or neighbors.
- Or, the penalty could be entirely internal, as when a person that violates his or her private rules of conduct anticipates feeling guilty afterwards.

iii. We now incorporate penalty F (a fine) into the game matrix.

- Clearly F>C will solve the dilemma, but will other smaller fines also work?
- How large does F have to be solve this public good problem?

A Public Policy Solution to a Public Good Game

		Bob	
		Povide	Free Ride
Al	Provide	(A, B) (2CP-C, 2CP-C)	(A, B) (CP-C, CP-F)
	Free Ride	CP-F, CP - C	(-F, -F)

- To solve the problem, CP-F < 2CP - C and CP-C > -F (Explain why.)
- Show that the first requirement is satisfied when F> C-CP or F> (1-P)C
- Show that the second requirement is satisfied when the same condition is met.
- The smallest possible fine is thus $F^{min} = (1-P)C + e$ where e is the smallest difference in payoffs that a player responds to (can observe).
- This implies that the smallest required fine falls with the productivity of contributions. [Why?]
- [Hint: Show this mathematically by differentiating F^{min} with respect to P]

iv. Similar matrices can be used to illustrate essential features of what economists call externality problems.

E. The above suggests that a wide variety of types of penalties can be used to solve PD dilemmas.

- Some will involve formal legal sanctions: the use of police and courts to impose fines.
- Others may be solved by the creation of organizations that can impose modest punishments, as when managers may reduce the salaries of their employees, deny future raises, or fire free riders (shirkers).
- Some such penalties will be only stochastically imposed, which will require using expected values to model their effectiveness and calculating minimal fines (as in Tullock's Logic of the Law or Becker's classic paper on crime.)
- Culture and norms may also solve PD problems through by associating conditional "approval" with contributing to a public good (cooperation) and/or "shame" or loss of status with free riding (defection)