1
Introduction

- The study of economics, or more generally the social sciences, is an attempt to put structure on the way people behave and make decisions, and then to use this structure to analyze questions pertinent to the study of societies. The process of putting structure on the way people behave means that we, the social scientists, will choose ways to model people as decision makers, Then, we will embed our models of behavior into more general models of society, in order to help us understand what we see, and more importantly, help us predict what we cannot yet see. The ultimate goal is to prescribe policy recommendations, to the private or public sector, based on the simplistic, yet rigorous analysis that ties our hands.

- Game theory is one such attempt at a rigorous mathematical modeling approach of conflicts and cooperation between *rational* agents. In this respect, *rational* needs to be well-defined. Following the tradition of mainstream decision theory and economics, rational behavior of an agent is defined to be choosing actions that maximize his utility (or some form of payoff) subject to the constraints that he faces. This is clearly a caricature of reality, but it is a useful starting point, and in some cases works surprisingly well.

- Before game theory was introduced into economics and the social sciences, the mainstream approach to economic analysis was that of *competitive behavior.* The premise of the competitive model is that each agent takes *as given* the environment he is in, and basically assumes that his actions have no effect on the external, or *exogenous* environment, but only directly on his utility. This may be a fine assumption about how my purchase of toothpaste will affect the price of toothpaste, but will not be a good assumption for the effect of British Petroleum taking the price of oil as given, when it is considering doubling its production of crude oil. More generally, when agents are *not negligible* to the environment that they are in, the assumption of competitive behavior–taking the environment to be constant– fails to hold.

- When the competitive assumption fails, we need to modify our treatment of agents, or "players", to reflect their understanding of how their actions effect their own payoffs, and how they effect the environment in which they are in. Thus, for example, Coca-Cola may have a good idea on how that its own marketing campaign will affect its demand, and the consequences on Pepsi. Pepsi, in return, understands that its response will affect its own profits and the profits of Coke. To maximize profits, each company has to make assumptions on, or *predict* the behavior of the other, knowing that the other is going through a similar thought processes.

- This is where game theory comes in as an important contribution that expands the tool-kit of economic analysis beyond the limitations of the competitive model. In particular, it allows for *strategic* behavior in which agents, or "players", understand how their actions affect the environment they are in, and can go through "strategic reasoning" in analyzing the likely outcome of the situation, or "game" that they are in. The objective of game theory is to develop a set of tools that allows us to make predictions about how agents will, or ought to behave in different *strategic environments.*

- Like most of the tools in modern economic analysis, or more generally– *rational choice theory*, game theory is based on a set of mathematical abstractions that try to capture the essence of these strategic environments.

Such an abstraction allows the social scientist a way to understand the role played by the assumptions that are made, and how these assumptions lead to different results. The objective is to be able to shed light on different social situations, and the behavior that results from them.

Example: Cournot Duopoly

Imagine there are two identical firms, call them 1 and 2, that produce saffron. Assume that there are no fixed costs to produce saffron, and let the variable cost to each firm i, of producing quantity q_i of saffron be given by the cost function,

$$c_i(q_i) = q_i^2 \text{ for } i \in \{1, 2\}$$

The assumption in a competitive environment is that each firm will take a market price, p, as given, and will believe that its behavior cannot influence the market price. Given this assumption, each firm will choose a quantity to maximize its profits. Thus, each (symmetric) firm faces a profit maximization problem as follows:

$$\max_{q_i} p \cdot q_i - q_i^2 ,$$

To solve this problem, the First-Order Condition (FOC) is[1]

$$p - 2q_i = 0$$

and therefore, this optimality condition results in the way that each firm will choose its supply as a function of the price to maximize profits as follows:

$$q_i(p) = \frac{p}{2}.$$

(Note that this should not be new to students of economics who know that in a perfectly competitive market, the marginal cost as a function of quantity will represent the supply curve of a firm. The marginal cost for each firm is the derivative

[1] When we use the first order condition to find a solution, we need to check that the second order condition is satisfied as well, and that there are no corner solutions. Throughout the text we will just use first order conditions since in the problems selected there will be a unique optimum.

of the cost function, $MC_i(q_i) = 2q_i$, and when we equate marginal cost to price we exactly get the supply function shown above.)

Let the demand for saffron be given by the following equation,

$$p(q) = 100 - q,$$

where q represents the total quantity demanded, and $p(q)$ is the resulting price. Alternatively, we can write the market demand as a function of price to be $q_D = 100 - p$.[2]

To find the *competitive equilibrium*, we need to find a price p such that the sum of both firms' production at that price equals the demand at that price. Algebraically this is trivial: we need to find a price that equates

$$q_1(p) + q_2(p) = q_D(p) \ ,$$

or,

$$2 \cdot \frac{p}{2} = 100 - p \, ,$$

where the left hand side is the total market supply (the sum of the individual firm supplies) and the right hand side is the inverse of the demand function. This yields $p = 50$, and each firm produces $q_i = 25$. For those who are used to see the competitive equilibrium from a graphical illustration of the supply and demand curves, we add up the two firms' supply curves "horizontally", and the equilibrium is where total supply intersects the demand curve.

In this competitive equilibrium each firm is indeed maximizing profits when it takes the price, $p = 50$, to be given, and each firm's profits are equal to $\pi_i = 50 \cdot 25 - 25^2 = 625$.

Now consider a more sophisticate firm 1that realizes its effect on the market price, and produces 24 instead of 25? The price will have to change for demand to equal supply, and indeed, the new price will be $p(q) = \$51$, where $q = 24 + 25 = 49$ instead of 50. The profits of firm 1 will now be

$$\pi_1 = 51 \cdot 24 - 24^2 = 648 > 625.$$

[2] Recall that economists "reverse" the demand function and write price as a function of quantity, rather than quantity as a function of price. This allows for some nice graphical analyses that we will also use.

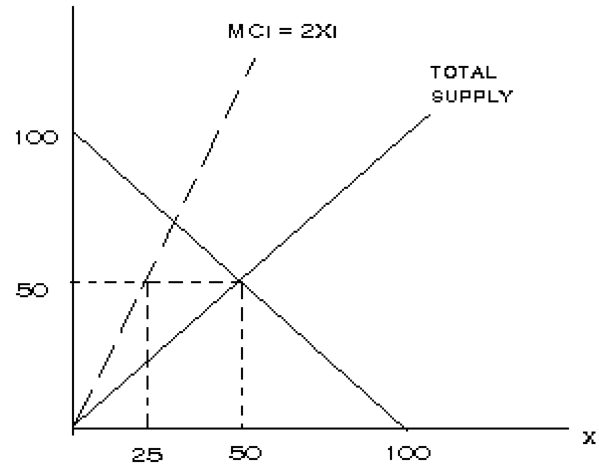

FIGURE 1.1.

Of course, once firm 1 realizes that it has such an effect on price, it should not just set $q_1 = 25$ but look for the best choice it can make. However, its best choice depends on the quantity that firm 2 will produce — what will that be? Clearly, firm 2 should be as sophisticated, and thus we have to find a solution that considers both *actions and counter-actions* of these rational and sophisticated firms.

- A variant of this example was studied first by Augustin Cournot in 1938, and his solution is a special case of the methods we will study. In what follows, we will study the methodology that has developed to address these strategic interactions, the methodology of *game theory.*

- The general theory has been applied in the social sciences to model interaction between consumers, firms, groups, political entities, etc., and has shed light on other disciplines such as evolutionary biology and psychology. However, the major focus of these notes, and the course they support, will be applied towards economics, with some sparse examples from political science, law, and life more generally.

- The notes are meant to be self contained. My aim to be concise and rigorous, yet I have added many examples to help the reader absorb the concepts of strategic analysis. Some mathematical maturity is useful to fully comprehend the material, yet it is not assumed that the reader has a strong mathematical background. Multivariate calculus, and knowledge of some basic probability theory are required to follow all of the material, but even without these some of the basic examples and constructions can be appreciated.

Part I

Static Games of Complete Information

2 Preliminaries

We are in search for a language that will help us think about strategic situations in which players who interact understand the environment they are in, and how their actions affect the outcomes that they and their counterparts will face. It is useful, therefore, to start with the simplest set of situations possible, and the simplest language that will capture these situations. A *static game of complete information* is the most fundamental game, or environment, in which such strategic considerations can be analyzed. It represents, for example, the Cournot duopoly set-up described earlier in the introduction.

The basic idea that a *static game* represents is a situation in which a set of players independently choose actions once-and-for-all, and given the actions that they chose some outcome will occur that affects each players well being. To be more precise, we can think of a static game as having two distinct steps:

Step 1: Each player *simultaneously and independently* chooses an action.

By *simultaneously and independently*, we mean something broader and more accommodating than players all choosing their actions at the exact same moment and without interacting to coordinate their actions. We mean that *information-wise*, players are making their choices without observing the choices that other players are making. For example, if you have to choose between two actions first, and I have to choose my action after you without observing your choice or knowing anything about it, then effectively it is *as if* we are choosing our actions simultaneously and independently.

Step 2: Conditional on the players' choices of actions, payoffs are distributed to each player.

That is, once the players all made their choices, these choices will result in a particular outcome, and this outcome generates some payoff, or utility, to each of the players in the game. For example, if we are playing rock-paper-scissors and I draw paper while you simultaneously draw scissors, then the outcome is that you win and I lose, and the payoffs are what winning and losing mean in our context (something physical like ten cents, or just the intrinsic joy of winning and suffering of losing that we so often experience).

Steps 1 and 2 above settle what we mean by *static*.What do we mean by *complete information?* The loose meaning is that *all players understand the environment they are in*, i.e., the game they are playing, in every way. This is not a trivial idea, as it turns out, because we need to give a more rigorous meaning to "understanding the environment".

Recall that we are seeking to depict a situation in which players can engage in *strategic reasoning*. That is, I want to predict how you will make your choice, given *my belief* that you understand the game. Your understanding incorporates *your belief* about my understanding, and so on. To make this sort of reasoning precise, we introduce the idea of *common knowledge*:

Definition 1 *An event E is* ***common knowledge*** *if (1) everyone knows E, (2) everyone knows that everyone knows E, and so on* ad infinitum*.*

On the face of it, this may seem like an innocuous assumption, and indeed, it may be in some cases. For example, if you and I are both walking in the rain together, then it is safe to assume that the event "it is raining" is common knowledge among us. However, if we are both sitting in class and the professor says "tomorrow there is an exam", then the event "there is an exam tomorrow" may not be common knowledge; even if I heard it, I cannot be sure that you did. Thus, requiring common knowledge is not as innocuous as it may seem, but without this assumption our life as analysts becomes very difficult.

Equipped with the notion of what *static* means, and what *common knowledge* means, we are ready to define the most fundamental of game representations:

Definition 2 *A* ***Static Game of Complete Information*** *is a static game in which the way players' actions affect each player's payoff is common knowledge to the players of the game.*

Going back to the duopoly game presented in the introduction, it would be considered a static game of complete information if the demand function and cost functions of the firms would be common knowledge to the two firms. Similarly, for rock-paper-scissors game that we play, it would be a static game of compete information if the rules, and what winning and losing meant in terms of payoffs, were common knowledge among us.

3

Representation of Games: The Normal Form

Now that we understand the basic ingredients of a static game of complete information, it will be useful to develop a language to represent it in a parsimonious, and rather general way, to capture the strategic essence of a game. The *normal-form* is one of the most common ways of representing a game, and it consists of the following three elements:[1]

1. A finite set of **players**, $N = \{1, 2, ..., n\}$
2. A (possibly infinite) set of **pure strategies** (or pure actions), S_i, for each players $i \in N$.
3. A set of payoff (or utility) functions for each player $u_i : S_1 \times S_2 \times ... \times S_n \to \Re$ for all $i \in N$.

This representation is very general, and will capture many situations in which each of the players $i \in N$ must *simultaneously* choose a feasible strategy (or ac-

[1]This is where we begin to use some mathematical notation that is convenient for summarizing formal statements. A finite set of elements will be written as $X = \{a, b, c, d\}$ where X is the set and a, b, c and d are the elements it includes. Writing $a \in X$ means "a is an element of the set A". If we have two sets, X and Y, we define the *Cartesian product* of these sets as $X \times Y$. If $a \in X$ and $h \in Y$ then we can write $(a, h) \in X \times Y$.

tion) $s_i \in S_i$. (We call a strategy *feasible* if it is a valid option.) After strategies are selected, each player will realize his or her payoff that will be given by $u_i(s_1, s_2, ..., s_n) \in \Re$ where $(s_1, s_2, ..., s_n)$ is the *strategy profile* (or action profile) that was selected by the agents, and $\Re$ represents the set of real numbers.[2]

Thus, the *normal form* will be a triple of sets: $\langle N, \{S_i\}_{i=1}^n, \{u_i\}_{i=1}^n \rangle$ where N is the set of players, $\{S_i\}_{i=1}^n$ is the set of all players' strategy sets and $\{u_i\}_{i=1}^n$ the set of all players' utility functions.[3]

Example: Prisoner's Dilemma

The Prisoner's Dilemma is probably the most well known example in game theory, and has some very general applications. It is a static game of complete information that represents a situation consisting of two individuals (players) who are suspects for a hideous crime. They are separated at the police station and questioned in different rooms, so each faces a choice between two actions: to say nothing to the investigators, also called "mum" (M), or to rat on the other suspect, also called "fink" (F). The payoff of each suspect is determined as follows: if both mum, then both get 1 year in prison since the evidence can support a minor offence; if, say, individual 1 mums while individual 2 finks, then individual 1 gets 9 years in prison while individual 2 gets to walk away free and vice versa if individual 1 finks while individual 2 mums; if both fink, then both get only 6 years because of conflicting arguments.

If we assume that each year in prison can be represented by -1 unit of utils, then we can represent this game in normal-form as follows:

- **players:** $N = \{1, 2\}$

- **strategy sets:** $S_i = \{M, F\}$ for $i \in \{1, 2\}$

[2] A function maps elements of one set (or a Cartesian product of sets) to elements of another. Here we encountered a utility function, that maps actions of all players into utils for each player. The set of all possible actions is the Cartesian product of the sets of actions for each player, $S_1 \times S_2 \times ... \times S_n$, and we will sometimes refer to it as S. A player's utility is given by a real number.

[3] $\{S_i\}_{i=1}^n$ is another way of writing $\{S_1, S_2, ..., S_n\}$. Similarly for $\{u_i\}_{i=1}^n$.

- **payoffs:** Let $u_i(s_1, s_2)$ be the payoff to player i if player 1 chooses s_1 and player 2 chooses s_2. We can then write payoffs as,

$$\begin{aligned} u_1(M,M) &= u_2(M,M) = -1 \\ u_1(F,F) &= u_2(F,F) = -6 \\ u_1(M,F) &= u_2(F,M) = -9 \\ u_1(F,M) &= u_2(M,F) = 0 \end{aligned}$$

This completes the normal form representation of the Prisoner's Dilemma.

Example: Cournot Duopoly

Taking the game we started with in the introduction, we have the following normal form representation:

- players are $N = \{1, 2\}$
- strategy sets are $S_i = [0, \infty]$ for $i \in \{1, 2\}$ and firms choose quantities $q_i \in S_i$
- payoffs are given by: for $i, j \in \{1, 2\}$, $i \neq j$,

$$u_i(x_{i,}x_j) = \begin{cases} (100 - q_i - q_j) \cdot q_i - q_i^2 & \text{if } q_i + q_j < 100 \\ -q_i^2 & \text{if } q_i + q_j \geq 100 \end{cases}$$

This completes the normal form representation of the Cournot Duopoly game. (Notice that here the payoff function is a little tricky because it has to be well defined for any pair of actions that the players choose. When we write down a demand function $p = 100 - q$ we are implicitly assuming that prices cannot fall below zero, so that if both firms produce a quantity that is greater than 100, we keep the price at zero, and the payoffs are just the firms' costs.)

Example: Voting on a new agenda

Imagine that there are three players that can remain in the status quo or vote in favor of a new agenda. For example, these can be three roommates who currently have an agreement that they clean the room once every two weeks (the status quo) and they are considering cleaning it every week (the new agenda). Each can vote "yes" (Y), "no" (N) or "abstain" (A). We can set all players' value from the status quo at 0. Players 1 and 2 prefer the new agenda and value it at 1, while player 3 dislikes the new agenda and values it at -1. The majority voting institution in place works as follows: if no votes are received, or if there is a tie, the status quo prevails. Otherwise, the majority is decisive.

We can represent this game in normal-form as follows:

- players: $N = \{1, 2, 3\}$
- strategy sets: $S_i = \{Y, N, A\}$ for $i \in \{1, 2, 3\}$
- payoffs: Let V be the set of action profiles for which the new agenda is chosen, and let Q be the set of action profiles for which the status quo remains. Formally,

$$V = \left\{ \begin{array}{ll} (Y,Y,N), & (Y,N,Y), \\ (Y,Y,A), & (Y,A,Y), \\ (Y,A,A), & (A,Y,A), \\ (Y,Y,Y), & (N,Y,Y), \\ (A,Y,Y), & (A,A,Y) \end{array} \right\} \text{ and } Q = \left\{ \begin{array}{llll} (N,N,N), & (N,N,Y), & (N,Y,N), & (Y,N,N), \\ (A,A,A), & (A,A,N), & (A,N,A), & (N,A,A), \\ (A,Y,N), & (A,N,Y), & (N,A,Y), & (Y,A,N), \\ (N,Y,A), & (Y,N,A), & (N,N,A), & (N,A,N), \\ (A,N,N) & & & \end{array} \right\}$$

Then, payoffs can be written as,

$$\begin{aligned} u_i(s_1, s_2, s_3) &= \begin{cases} 1 & \text{if } (s_1, s_2, s_3) \in V \\ 0 & \text{if } (s_1, s_2, s_3) \in Q \end{cases} \quad \text{for } i \in \{1, 2\}, \\ u_3(s_1, s_2, s_3) &= \begin{cases} -1 & \text{if } (s_1, s_2, s_3) \in V \\ 0 & \text{if } (s_1, s_2, s_3) \in Q \end{cases} \end{aligned}$$

This completes the normal form representation of the voting game.

As you can see from the voting game above, it will sometimes be quite cumbersome to write down the formal representation of a game, but doing this totally specifies what is at hand: who the players are, what they can do, and how their actions affect each and every player. It turns out that for two person games in which each player has a finite number of (pure) strategies, there is a convenient representation.

3.1 Matrix Representation: Two-Player Finite Game

In many cases, players may be constrained to choose one of a finite number of actions. This is the case for the prisoner's dilemma, rock-paper-scissors, the voting game described above and many more strategic situations. In fact, even when players have infinitely many actions to choose from, we may be able to provide an ample approximation by restricting attention to a finite number of actions. If we think of the Cournot duopoly example and we consider the production of cars, then first we can safely assume that we are limited to units of complete cars (which reduces the action set to the natural numbers, a countable infinity), and if we assume that flooding the market with more than 2 billion cars will cause the price of cars to drop to zero, then we have effectively restricted the strategy set to a finite number of strategies (two billion, to be accurate).

Being able to distinguish games with finite acti9on sets is useful, so we define,

Definition 3 *A **finite game** is a game with a finite number of players, in which the number of strategies (actions) in S_i is finite for all $i \in N$.*

As it turns out, any two-player finite game can be represented by a matrix in the following way:

1. **Rows:** each row represents one of player 1's strategies

2. **Columns:** each column represents one of player 2's strategies

3. **Entries:** Each entry in this matrix contains a 2-element vector (u_1, u_2), where u_i is player i's payoff when the actions of both players correspond to the row and column of that entry.

Example: The Prisoner's Dilemma in Matrix Form

Recall that in the Prisoner's Dilemma each player had two actions, M (mum) and F (fink). Therefore, our matrix will have two rows (for player 1) and two columns (for player 2). Using the payoffs for the prisoner's dilemma given in the example above, the matrix representation of the Prisoner's Dilemma is:

Player 2

		M	F
Player 1	M	-1,-1	-9,0
	F	0,-9	-6,-6

Arguably, this is a much simpler way of representing the game, and all the information appears in a concise and clear way. Note, however, that neither of the other two examples given above (Cournot and voting) can be represented by a matrix: the Cournot Duopoly game is not a finite game, and the voting game has more than two players.[4]

(**Note**: such matrices are often referred to as *bi-matrices*. he reason is that in a traditional matrix, by definition, each entry corresponding to a row-column combination must be a single number, or element, while here each entry has a vector of two elements – the payoffs for each of the two players. Thus, we formally have *two matrices*, one for each player. We will adopt the common abuse and call this a matrix.) L2

[4] We can represent the voting game using three 3×3 matrices: the rows of each matrix represent the actions of player 1, the columns those of player 2, and each matrix corresponds to an action of player 3. However, the convenient features of two player matrix games are harder to use for three player, multiple matrix representations, not to mention the rather cumbersome structure of multiple matrices.

4
Solution Concepts

Now that we have found a concise language in which to represent games, we would like to go a step further and be able to make some prediction, or be able to prescribe behavior, about how players should (or will) play.

Let's begin with the Prisoner's Dilemma given above, and imagine that you were to advise player 1 about how to behave (you're his lawyer). Being a thorough person, you make the following observation for player 1: "If player 2 chooses F, then playing F gives you -6, while playing M gives you -9, so F is better." Player 1 will then bark at you and say "my buddy will never squeal on me!" You, however, being a diligent lawyer, must reply with the following answer: "If player 2 chooses M, then playing F gives you 0, while playing M gives you -1, so F is better. It seems like F is always better!"

Indeed, if I were the lawyer of player 2 then the same analysis works for him, and this is the "dilemma": each player is better off playing F regardless of his opponents actions, but this leads them both to receive payoffs of -6, while if they can only agree to both choose M, then they would obtain -1 each. Left to their own devices, and to the advocacy of their lawyers, the players should not be able to resist the temptation to choose F. Even if player 1 believes that player 2 will play M, he is better off finking (and vice versa).

Your intuition might want to convince you otherwise. You might want to say that they are friends, stealing together for some time now and care for each other. In this case, one of our assumptions is incorrect: the numbers in the matrix may not represent their true utilities, and if taken into consideration, altruism would lead to M chosen by both.

For example, to capture the idea of altruism and mutual caring, we can assume that a year in prison for each player is worth -1 for himself, and imposes another pain worth $-\frac{1}{2}$ to the other player. In this case, if player 1 chooses F and player 2 chooses M then player 1 gets $-4\frac{1}{2}$ ($-\frac{1}{2}$ for each year player 2 goes to jail) and player 2 gets -9 (only he spends time in jail). The actual matrix would then be,

		Player 2	
		M	F
Player 1	M	$-1\frac{1}{2},-1\frac{1}{2}$	$-9,-4\frac{1}{2}$
	F	$-4\frac{1}{2},-9$	$-9,-9$

and always playing M is as good as it gets. This shows us that out results will, as they always do, depend crucially on our assumptions.[1]

Another classic example is *the battle of the sexes* (Duncan Luce and Howard Raiffa, *Games and Decisions*, Wiley 1957). The story goes as follows. Pat and Chris are a couple and they need to choose, while each is at work and they have no means of communicating, where to meet this evening. Both prefer being together over not being together, but Pat prefers Opera (O) to Football (F), while Chris prefers the opposite. This is summarized in the following matrix:

		Chris	
		O	F
Pat	O	2,1	0,0
	F	0,0	1,2

What can you recommend to each now? Unlike the Prisoner's Dilemma, the best action for Pat depends on what Chris will do and vice versa. If we want to predict,

[1] Another change in assumptions might be that player 2's brother is a psychopath, and if player 1 finks then 2's brother will kill him, giving player 1 a utility of, say, $-\infty$ from finking.

or prescribe actions for this game, we need to make assumptions on *behavior* and on *beliefs* of the players. We are now in search of a *solution concept*, that will result in predictions.

We will define a solution concept as a way of analyzing games so as to restrict the set of all possible outcomes to those that are more reasonable than others. That is, we will look for some reasonable assumptions on behavior and beliefs of players that will hopefully divide the space of outcomes into "more likely" and "less likely". Furthermore, we would like our solution concept to apply to a large set of games. First, we define a solution concept more formally,

Definition 4 *A* ***solution concept*** *is a mapping from games to actual strategies, i.e.,*

$$F : \langle N, \{S_i\}_{i=1}^n, \{u_i\}_{i=1}^n \rangle \rightrightarrows S$$

where $S = S_1 \times S_2 \times ... \times S_n$ *is the set of all strategy profiles. That is, a solution concept* F *will prescribe a* ***subset of strategy profiles*** $S' \subset S$ *as its prediction.*

We will often use the term *equilibrium* for any one of the strategy profiles that emerge as one of the solution concept's predictions. As alluded to above, we will often think of equilibrium as the *actual predictions* of our theory. A more forgiving meaning would be that equilibria are the *likely predictions* since our theory will often not account for all that is going on, and in some cases we will see that more than one equilibrium prediction is possible. In fact, this will sometimes be a strength, and not a weakness of the theory.

4.1 Assumptions and Set-Up

To set up the background for equilibrium analysis, it is useful to revisit the assumptions that we will be making throughout:

1. Players are "rational"

   A *rational* player is one who chooses his action, $s_i \in S_i$ to maximize his payoff consistent with his beliefs about what others are doing.

2. Players are *"intelligent"*

   An *intelligent* player knows as much about the game as the game theorist.[2]

3. The game (which includes 1 and 2 above) is *common knowledge* among the players

Adding to these three assumptions, which we discussed a bit earlier, is a fourth assumption that constrains the set of outcomes that are reasonable:

4. Any prediction (or equilibrium) must be *self-enforcing*

   No third party that can enforce any profile of actions that the players will not voluntarily agree to stick to

The requirement that any equilibrium must be self-enforcing is at the core of our analysis in that we will assume that the players engage in *Noncooperative Behavior*.[3] That is, if a profile is to be an equilibrium, we will require each player to be happy with his choice, if he or she correctly predicted what the other players are doing in that profile. As one can by now figure out, the profile (F, F) is self-enforcing in the Prisoner's Dilemma game: each player is happy playing F.

The requirement of self enforcing predictions is a natural one if we take the game to describe the environment. If there are outside parties that can, through force, enforce profiles of strategies, then this is not represented by the payoffs that our game lays out. In such a case we ought to model the third party as a player, who has actions (strategies) that describe the enforcement.

4.2 Evaluating Solution Concepts

By developing a theory that will ultimately prescribe some predictions about the behavior of players in games, we want to be able to evaluate our theory as a

[2] This is not assumed in competitive equilibrium theory: if players can change prices, then game theory assumes that they know this, while competitive equilibrium theory does not.

[3] This is one of the fundamental ways in which non-cooperative game theory differs from cooperative game theory. For more on cooperative game theory see ...

methodological tool. That is, an method that should work for analyzing games that describe the strategic situations we are interested in. We will introduce four criteria that will help us evaluate a variety of solution concepts.

I. Existence: how often does it apply?

A solution concept is valuable insofar as it applies to a wide variety of games, and not just to a small and select family of games. One would hope that the solution concept is general in applications and not developed in an ad-hoc way that is specific to a certain situation or game. That is, when we apply our solution concept to a game we would hope that it will result in the *existence* of a solution.

For example, if a concept has the following prescription: "when players look each-other in the eye, then they will choose the best outcome among all those that have equal payoffs to all players." If this is our theory, our "solution concept", then it may fail to apply to many–maybe most–strategic situations. When players either don't "look each-other in the eye", or when they do not have equal payoff outcomes, then this theory would be silent about which outcomes are more likely to emerge as equilibrium outcomes.

As you can imagine, any proposed theory of a solution concept that relies on very specific elements of a game will not be general in that it will be hard to adapt it to a wide variety of strategic situations. Any such construction would make the proposed theory quite useless beyond the very special situations it was tailored to address. Thus, one goal is to have a method that will be general enough and apply to many strategic situations, that is, it will prescribe a solution that will *exist* for most games we can think of.

II. Uniqueness: How much does it add?

Just as we want our solution concept to offer some guidance for any game, we would like that guidance to be meaningful in that it restricts the set of possible outcomes, to a smaller set of reasonable outcomes. In fact, one may argue that being able to pin-point a single, *unique* outcome as a prediction would be ideal. *Uniqueness* is then an important counterpart to *existence*.

For example, if the proposed solution concept says "anything can happen," it always exists – regardless of the game we apply this concept to, "anything can happen" will always say that the solution is one of the (sometimes too many) possible outcomes. It goes without saying that this solution concept is worthless. A good solution concept is one that balances existence, so that it works for many games, with uniqueness, so that we can add some intelligent insight above the triviality of what is possible can happen.

III. Invariance: How robust is it?

The discussion of existence and uniqueness as desirable properties for a solution concept was not difficult to make, and hopefully, it was not difficult to follow. A third criterion that is somewhat more subtle, is important in qualifying a solution concept as a reasonable one, and this criterion requires that a proposed solution concept should be *invariant* to small changes in the game's structure.

What I mean by "small changes" needs to be qualified more precisely. Adding a player to a game, for instance, may not be a small change if that player has actions that can wildly change the outcomes of the game without that player. Thus, changes in the set of players cannot innocuously be considered small changes in the game. Similarly, if we add or delete strategies from the set of actions that are available to a player, we may hinder his ability to guarantee himself some outcomes, which again should not be considered a small change to the game. We are left only with one component to fiddle with: the payoff (or utility) functions of the players. It is reasonable to argue that if the payoffs of a game are modified trivially, then this would be a small change to the game that should not affect the predictions of a "robust" solution concept.

For example, consider the Prisoner's Dilemma game we introduced earlier. If instead of 9 years in prison imposing a pain of -9 for the players, it imposed a pain of -9.01 for player 1 and -8.99 for player 2, we should be somewhat discouraged if our solution concept suddenly changes the prediction of what players will, or ought to do.

Thus, *invariance* is a kind of robustness property that we would like a solution concept to comply with. In other words, if two games are "close" so that the action sets and players are the same yet the payoffs are trivially different, then our solution concept should offer predictions that a wildly different for the two games.

4.3 Evaluating Solutions

We have laid out criteria to evaluate the methodological value of a solution concept in way that gives us guidance with respect to its desirable properties. Once we ascribe to any particular solution concept, as social scientists we would like to evaluate the properties of the *solutions,* or predictions, that our solution concept will generate and prescribe. This is useful insofar as to give us insights into what we expect the players of a game to achieve if they are left to their own device, and in turn, it can give us guidance into ways of possibly changing the environment to improve upon the social outcomes of the game.

Once more, we have to be precise about what it means to "improve upon" the social outcome. For example, many people may have agreed that in 1910 it would be socially better to take away ten dollars from the very rich Henry Ford, and give those ten dollars to an assembly line worker in his Model-T production plant. In fact, maybe even Henry Ford would have approved of this transfer if, say the worker's child was ill, and those ten dollars would have saved the child's life. However, Henry Ford may or may not have liked that idea, especially if such government intervention would imply that in the future most of his wealth would be dissipated with such transfers.

There is an important, and non-controversial criterion that economists use for evaluating whether an outcome is socially desirable. This is the criterion of *Pareto Optimality.* We say that an outcome is Pareto optimal (or Pareto efficient) if we cannot find another outcome that makes some people better off, and no people worse off. The criterion of Pareto optimality is in tune with the idea of efficiency, or "no waste", in that all the possible value of the interaction is in some way distributed among the players. This should indeed be non-controversial since we are evaluating people's payoffs not as a function of the monetary value, but as a function of their *true payoffs*, or utility.

To make the point clear, imagine that players are altruistic, or, for example, care about relative wealth. Then, to evaluate the story of Henry Ford's transfer above, we need to translate monetary distribution outcomes into a notion of payoffs, and evaluate outcomes with the *payoffs they generate* and not the distribution of wealth. If, for example, I would rather have you and I get five dollars each over me getting ten dollars, then my *payoff* from the outcome "both of us get \$5" must be greater than my payoff from the outcome "I get \$10 and you get nothing". When we write down the utility of a player from an outcome in a game, we will must always take into account the payoff from the outcome, and not just monetary transfers so that our analysis is accurate and correct.

We now introduce the notion of Pareto optimality formally:

Definition 5 *A strategy profile $s \in S$* ***Pareto dominates*** *strategy profile $s' \in S$ if $u_i(s) \geq u_i(s') \forall\ i \in N$ and $u_i(s) > u_i(s')$ for at least one $i \in N$ (in which case, we will also say that s' is* ***Pareto Dominated*** *by s.) A strategy profile is* ***Pareto optimal*** *if it is not Pareto dominated by any other strategy profile.*

As social scientists we will hope that players will find ways to coordinate on Pareto optimal outcomes, and will avoid those that are Pareto dominated. However, as we will see time and time again, this will not be the case in many games. For example, in the Prisoner's Dilemma we made the case that (F, F) should be considered as a very likely outcome, in fact, as we will argue forcefully later, *the* likely outcome. Once can see, however, that it is Pareto dominated by (M, M). (Notice that $(M, M), (M, F)$ and (F, M) are all Pareto optimal since no other profile dominates any of them. Don't confuse Pareto optimality with the best "symmetric" outcome that leave all players "equally" happy.)

We are now ready turn to some specific solution concepts, and see how they fair with respect to the criteria described above.

5

Strict Dominance in Pure Strategies

Our prescription to the Prisoner's dilemma was easy: they both had an action that was best *regardless* of what their opponent chose. The search for such a prescription can be generalized as follows:

Definition 6 *Let $s_i \in S_i$ and $s_i' \in S_i$ be feasible strategies for player i. We say that s_i' is* ***strictly dominated*** *by s_i if for any feasible combination of the other players' strategies, player i's payoff from s_i' is strictly less than that from s_i. That is,*

$$u_i(s_1, s_2, ..., .s_i, ...s_n) > u_i(s_1, s_2, ..., s_i', ...s_n)$$

for all $(s_1, s_2, ..., s_{i-1}, s_{i+1}, ...s_n) \in \underbrace{S_1 \times S_2... \times S_{i-1} \times S_{i+1}... \times S_n}_{n-1 \text{ strategy spaces}}$. *We will write $s_i \succ_i s_i'$ to denote that s_i' is strictly dominated by s_i .*

For example, recall that in the Prisoner's dilemma, playing M was worse than playing F regardless of what one's opponent did. Thus, M is strictly dominated by F for both players. In the Cournot duopoly game that we saw in the introduction, any quantity $x_i > 100$ will yield negative profits (price is zero, and production costs are positive) so any such quantity is strictly dominated by $x_i = 0$, which guarantees zero profits.

It is quite evident that the formal notation used for the definition of a dominant strategy is somewhat long-winded when we have to specifically write down the strategies, or actions chosen by the players who are not player i, namely the profile

$$(s_1, s_2, ..., s_{i-1}, s_{i+1}, ...s_n) \in S_1 \times S_2 ... \times S_{i-1} \times S_{i+1} ... \times S_n .$$

For this reason it will be very useful, and time saving, to use a common shorthand notation as follows: We define $S_{-i} \equiv S_1 \times S_2 ... \times S_{i-1} \times S_{i+1} ... \times S_n$ as the set of all the strategy spaces of all players who are not player i. We then define $s_{-i} \in S_{-i}$ as a particular feasible profile of strategies for all players who are not i. Then, we can rewrite the formal inequality of s_i' being strictly dominated by s_i as,

$$u_i(s_i, s_{-i}) > u_i(s_i', s_{-i}) \text{ for all } s_{-i} \in S_{-i}.$$

What this says is that *given any* strategy of i's opponents, player i gets more utility from playing s_i than from playing s_i'.

Since a strictly dominant strategy is one to avoid at all cost,[1] there is a counterpart strategy that would be very desirable, if you can find it. This is strategy that is always the best thing you can do, regardless of what else is going on in the game. Formally:

Definition 7 *$s_i \in S_i$ is a **strictly dominant strategy** for i if every other strategy of i is strictly dominated by it, that is,*

$$u_i(s_i, s_{-i}) > u_i(s_i', s_{-i}) \textit{ for all } s_i' \in S_i \textit{ and } s_{-i} \in S_{-i}.$$

Once again, we can refer to the Prisoner's dilemma as an example. As we argued earlier, each of the two players had a strategy that was his best choice *regardless of his opponents behavior*, namely, to fink. Therefore, in the Prisoner's dilemma, each player has a dominant strategy: fink!

If, as in the Prisoner's dilemma, every player had such a wonderful dominant strategy, then this would be a very sensible prediction for behavior. That is, if each player has an action that is his best choice regardless of the other players' choices,

[1] This is a good point to stop and reflect on a very simple, yet powerful lesson. When you are in any situation, look first for your dominated strategies and avoid them!

we should expect the game to result in one particular outcome: that the players all choose their dominant strategy. More generally, we can introduce this simple idea as our first solution concept:

Definition 8 *The strategy profile $s^D \in S$ is a* ***strict dominant strategy equilibrium*** *if for all $i \in N$, $s_i^D \in S_i$ is a strict dominant strategy.*

Now we have a formal name for the outcome "both player fink", or (F, F), in the prisoner's dilemma: it is a dominant strategy equilibrium. If we apply this concept to any game, it basically says that we should identify the strict dominant strategies of each player, and then use this profile of strategies to predict or prescribe behavior. As it turns out, when we can do this, like we did in the Prisoner's dilemma, this concept has a very appealing property:

Proposition 1 *If the game $\Gamma = \langle N, \{S_i\}_{i=1}^n, \{u_i\}_{i=1}^n \rangle$ has a strictly dominant strategy equilibrium s^D, then s^D is the* ***unique*** *dominant strategy equilibrium.*

This is rather easy to prove, and is left for you as an exercise. The proposition is very useful in giving an answer about one of our criteria to evaluate a solution concept: this solution concept guarantees uniqueness. However, what do we know about existence? A quick observation will easily convince you that this is a problem.

Consider, for example, the Battle of the Sexes game introduced earlier. Neither player has a dominant strategy, and since each player has only two strategies, neither has a dominated strategy either. The best strategy for Chris depends on what Pat is doing, and vice-versa. Thus, if we stick to the solution concept of strict dominance we will encounter games, in fact many of them, for which there will be no equilibrium, and in turn, will give us no insight into the choices that player ought to, or will make.

Going back to the Prisoner's dilemma, where we did have an equilibrium solution using the strict dominance solution concept, we can evaluate the efficiency properties of the unique strictly dominant strategy equilibrium. It is easy to see that the outcome prescribed by this solution is not Pareto optimal: both players would be better off if they can commit to play "mum", yet left to their own incentives, they will not do this.

The failure of Pareto optimality is *not a* failure of the solution concept. Given the payoffs, and the assumption that players are rational in that they seek to maximize their payoffs, this is exactly what we would expect them to do. The failure of Pareto optimality implies that the players would benefit from modifying the environment they are in. For example, if a Mafia boss can enforce implicit agreements so that people who fink on someone are hurt severely, then this will change the payoffs to the game and in turn may lead to an equilibrium in which both players are better off.

For instance, if the two prisoners could join the Mafia, and if the code of conduct is that when someone finks on a member of the Mafia then he is reprimanded severely (very severely!), then this will change the payoff structure of the Prisoner's dilemma if they are caught. In particular, imagine that the pain from Mafia punishment was equivalent to -10, in which case we have to subtract 10 units of utility for each player who finks and the Mafia-modified Prisoner's dilemma is represented by the following matrix:

		Player 2	
		M	F
Player 1	M	$-1,-1$	$-9,-10$
	F	$-10,-9$	$-16,-16$

in which case M becomes the strict dominant strategy, and the prediction of the strict dominance solution is that (M, M) will be the unique outcome of the game.

Remark 1 *We will sometimes use the notion of* ***weak dominance****. We say that s_i' is* weakly dominated *by s_i if for any feasible combination of the other players' strategies, player i's payoff from s_i' is weakly less than that from s_i, That is,*

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \text{ for all } s_{-i} \in S_{-i}.$$

This means that for some s_{-i} this inequality can strictly hold, while for others it will hold with equality. We will define a strategy to be ***weakly dominant*** *in a similar way. This is still useful since if we can find a dominant strategy for a player, be it weak or strict, this seems like the most obvious thing to prescribe.* L3

6

Iterated Elimination of Strictly Dominated Strategies

With strict dominance we asked whether a player had a strategy that is strictly dominant, or *always best*, and if he had such a strategy we prescribed that as the behavior we should expect in equilibrium. We saw, however, that such a strong solution will often fail to exist. This implies that a general theory will have to be less demanding than that of choosing strictly dominant strategies. In this section we consider a less extreme concept that is based on the idea of dominance, and common knowledge of rationality, but does not require the existence of a strictly dominant strategy.

The concept we introduce here is called *Iterated Elimination of Strictly Dominated Strategies* (IESDS), and builds on the premise that rational players will not play strictly dominated strategies. Using our central assumptions that both the payoffs of the game and the rationality of the players is common knowledge, this simple observation gives us a rather simple and appealing procedure that will give us guidance in predicting what behavior *should not* prevail, which in turn may reduce the set of predicted outcomes.

More specifically, if all the players know that each player will never play a strictly dominated strategy, they can "erase" the possibility that dominated strategies will be played by any player. This step may reduce the *effective* strategy space for each of the players, thus defining a new "smaller" game. But, since it is common knowledge that in this reduced game players will not play strictly dominated strategies, this process can be repeated for the reduced game. Indeed there may be strategies that were not dominated in the original game, but are in the new reduced game. Since it is common knowledge that players will perform this kind of reasoning again, the process can continue until no more strategies can be eliminated in this way. We can think of this process as an algorithm that works as follows:

Step 1: Consider the strategy sets of each player.

Step 2: For each player, are there strategies that are strictly dominated?

Step 3: If the answer in 2 is 'yes', eliminate these strategies and define a new game, go back to step 1. If the answer is 'no' for all players then go to step 4.

Step 4: The remaining strategies are reasonable predictions for behavior.

We will call any strategy profile $s^* = (s_1^*, ..., s_n^*)$ that survives this process of IESDS, an *iterated-elimination equilibrium*.

Remark 2 *In this chapter we will refrain from giving a precise mathematical definition of the process, since to do this we need to consider richer behavior by the players, which includes their ability to randomly choose between their different pure strategies, where a pure strategy is any $s_i \in S_i$. We will revisit this briefly when such randomizations, or* mixed strategies *are introduced later. Just to satisfy your curiosity, think of the battle of the sexes, and imagine that Chris can pull out a coin and flip between the decision of opera or football – this by itself introduces a "new" strategy!*

A Simple Example

Consider the following two-player discrete game:

	L	C	R
U	4,3	5,1	6,2
M	2,1	8,4	3,6
D	3,0	9,6	2,8

Notice first that there is no dominant strategy for player 1 or for player 2. Also note that there is no strictly dominated strategy for player 1. There is, however, a strictly dominated strategy for player 2: C is strictly dominated by R. Thus, since this is common knowledge, both players know that effectively we can eliminate C, which results in the following reduced game:

	L	R
U	4,3	6,2
M	2,1	3,6
D	3,0	2,8

In this reduced game both M and D are strictly dominated by U for player 1, which could only be concluded after the first stage of reasoning we did above. Eliminating these two strategies yields the following trivial game:

	L	R
U	**4,3**	6,2

in which player 2 has a strictly dominant strategy, playing L. Thus, for this example IESDS yielded a unique prediction: the strategy profile we expect these players to play is (U, L), yielding the payoffs of (4,3).

Remark 3 *Be careful, the pair of payoffs (4,3) is* ***not a strategy****. Strategies are a set of actions by the players, and payoffs are a result of the outcome. When we will talk about predictions, or equilibria, we will always refer to what players do as the equilibrium, not their payoffs.*

A More Complicated Example: Cournot Revisited

Consider the following example of a Cournot Duopoly which is similar to the example we introduced earlier, but is simpler in that the firms have linear rather than quadratic costs:

- Cost for firm i, $c_i(q_i) = c_i \cdot q_i$ for $i \in \{1, 2\}$
- Demand: $p(q) = 100 - q$ where $q = q_1 + q_2$

Let's make things even simpler and assume that marginal costs (and total costs) of production are zero, i.e., $c_i = 0$ for $i \in \{1, 2\}$.[1]

Now consider the profit function of firm 1. Once firm 1 understands the game it is playing, it must realize that the profit function it faces is the following:

$$\pi(q_1, q_2) = p(q) \cdot q_1 = (100 - q_1 - q_2) \cdot q_1 .$$

What should firm 1 do? If it knew what quantity firm 2 will choose to produce, say some value of q_2, then the profits of firm 1 are maximized when the first order condition is satisfied:[2]

$$q_1 = \frac{100 - q_2}{2} . \tag{6.1}$$

Though it is true that the choice of firm 1 depends on what it thinks form 2 is doing, equation (6.1) implies that firm 1 will *never* choose to produce more than 50. This follows from the simple observation that $q_2 \geq 0$. If you prefer a graphical representation, this conclusion follows from the shape of the profit function: for any value of q_2, firm 1's profits from choosing $q_1 = 50$ are higher than its profits from choosing a quantity above 50.

This simple observation leads to our first step of iterated elimination: a rational firm 1 produces no more than 50, implying that the effective strategy space is $q_1 \in [0, 50]$. Clearly, a similar argument works for firm 2. Using this first stage of elimination, we can turn to the second stage of elimination. Since $q_2 \leq 50$, equation

[1] This zero cost assumption may strike you as crazy, but it would not matter to the qualitative analysis if we add positive marginal costs. In fact, in the current day and age of software packages, the zero cost assumption of production (not research and development) is not that crazy!

[2] This is a necessary, not sufficient condition. As you know, one has to consider the second order condition, but in this case things work out just fine.

(6.1) implies that firm 1 will choose a quantity no less than 25, and the second stage of elimination implies that $q_i \in [25, 50]$ for $i \in \{1, 2\}$.

You should be able to convince yourself that the next stage will reduce production to $q_i \in [25, 37\frac{1}{2}]$, and this process can continue on and on. Interestingly, the set of strategies that survives this process converges to a single quantity choice of $q_i = 33\frac{1}{3}$.[3]

We can now turn to evaluate the IESDS solution concept using the criteria we introduced earlier. Let's start with existence. Unlike the concept of strict dominance, we can apply IESDS to *any* game. All we need to do is apply the algorithm described above, which does not require the existence of a strictly dominant strategy, nor does it require the existence of strictly dominated strategies. It is the latter, however that gives this concept some bite: when strictly dominated strategies exist, the process of IESDS is able to say something about how rationality restricts behavior.

It is nice to know that unlike strict dominance, an IESDS solution always exists. The fact that a set of strategies that survives IESDS always exists is at the cost of uniqueness, however. In the two examples above, IESDS lead to the survival of a unique strategy. However, consider the battle of the sexes:

		Chris	
		O	F
Pat	O	2,1	0,0
	F	0,0	1,2

IESDS is not useful here: neither O nor F are strictly dominated strategies for the players. As we can see, this solution concept can be applied (it exists) to any game, but it will often fail to provide a unique solution. We now see that uniqueness, or very tight predictions, are at odds with existence, or wide applicability. This must

[3] If this were to converge to an interval, then by the symmetry between both firms, the resulting interval for each firm would be $[q_{\min}, q_{\max}]$ and two equations would hold: $q_{\min} = \frac{100 - q_{\max}}{2}$ and $q_{\max} = \frac{100 - q_{\min}}{2}$. The only solution is thus $q_{\min} = q_{\max} = 33\frac{1}{3}$.

be taken into account by the scholar who is in search of a solution concept: it should be widely applicable, but be able to have bite on making some restricted, and hopefully falsifiable predictions.

The social desirability of IESDS equilibria should have been anticipated. An easy example is the Prisoner's dilemma. There, IESDS quickly leaves (F, F) as the unique survivor, or equilibrium, and it is not Pareto optimal. Similarly, both examples above (the 3×3 matrix and the Cournot game) provide further evidence that Pareto optimality need not be achieved by IESDS: In the 3×3 matrix example, both strategy profiles (M, C) and (D, C) yield higher payoffs for both players, $(8, 4)$ and $(9, 6)$ respectively, than the unique IESDS equilibrium, which yields $(4, 3)$. For the Cournot game, producing $33\frac{1}{3}$ each yields profits of $1, 111\frac{1}{9}$ for each firm. If instead they would both produce 25 (and together gain monopoly profits this way) then each would earn profits of $1, 250$. Thus, common knowledge of rationality, does not mean that players can guarantee the best outcome for themselves when their own incentives dictate their behavior.

On a separate note, it is interesting to observe that there is a very simple, and quite obvious relationship between the two concepts we have seen so far:

Proposition 2 *If for a game* $\Gamma = \langle N, \{S_i\}_{i=1}^n, \{u_i\}_{i=1}^n \rangle$, s^* *is a* ***strict dominant strategy equilibrium****, then* s^* ***uniquely survives IESDS****.*

This proposition is very easy to prove: If $s^* = (s_1^*, ...s_n^*)$ is a strict dominant strategy equilibrium then by definition for every player i, all other strategies s_i' are strictly dominated by s_i^*. This implies that after one stage of elimination we will be left with a single profile of strategies, which is exactly s^*, and this concludes the proof.

Notice that the ideas behind both of the concepts we have seen are based on eliminating actions that players would never take. An alternative approach would be to ask what players might choose to do, and under what conditions, which is the focus of the next section.

7

Beliefs, Best Responses and Rationalizability

When we considered eliminating strategies that no rational player would choose to play, it was by finding some strategy that is always better, or as we said, dominates the eliminated strategies. What is special about a strategy that cannot be eliminated? It means that under some conditions, there is no better strategy, or in other words, under some conditions, this strategy is the *best response* the player can choose.

When we qualify a strategy to be the best one *under some conditions*, these conditions must be expressed in terms that are rigorous, and are related to the game played. For example, think about situations when you were puzzled about the behavior of some individual given the situation he was in. To qualify his choice as irrational, or simply stupid, you would have to consider whether there is a way in which he can defend his action as a good choice. One natural way to get at this is to simply ask him: "what were you thinking?" If the response lays out a plausible situation for which his choice was a good one, then you cannot question his rationality.

This is where we are formally going in this section. If a strategy, say s_i is not strictly dominated for player i then it must be that there are combinations strategies of player i's opponents for which the strategy s_i is player i's best choice. This reasoning will allow us to *rationalize* the choice of player i.

7.1 The Best Response

The idea of choosing a best strategy as a response to the actions of ones opponents is captured with the following formal definition:

Definition 9 *$s_i \in S_i$ is a* ***best response*** *(BR) to $s_{-i} \in S_{-i}$ if*

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \; \forall \, s_i' \in S_i.$$

This definition will escort us from now until the end of the text. It is central to the notion of strategic behavior and rationality: a rational player, faced with a certain scenario, will always choose a best response to that scenario. For instance, take the Battle of the Sexes:

		Chris	
		O	F
Pat	O	2,1	0,0
	F	0,0	1,2

If Chris was convinced that Pat will go to the opera, then Chris's best response is to go to the opera because

$$u_2(O, O) = 1 > 0 = u_2(O, F) \; .$$

Similarly, if Chris believes that Pat will go to the Fight, then Chris's best response is to go to the fight.

There are some simple and nice relationships between the notion of playing a best response, and the notion of dominated strategies. First, if a strategy s_i is strictly dominated, it means that some other strategy s_i' is *always* better. But then, the strategy s_i could not be a best response to *anything*. This simple and intuitive argument leads to the following formal proposition:

Proposition 3 *If s_i is a* ***strictly dominated strategy*** *for player i, then it* ***cannot be a best response*** *to any s_{-i}.*

proof: If s_i is strictly dominated, then there exists some $\widehat{s}_i \succ_i s_i$ such that $u_i(\widehat{s}_i, s_{-i}) > u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$. But this in turn implies that there is no $s_{-i} \in S_{-i}$ for which $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})\ \forall\, s_i' \in S_i$, and thus s_i cannot be a best response to any $s_{-i} \in S_{-i}$. ■

An extreme companion to the proposition above would explore strictly dominant strategies. You should easily be able to convince yourself that if a strategy s_i is a strictly dominant strategy, then is is a best response to *everything* i's opponents can do. This immediately implies the next proposition, which is a direct consequence of this simple intuition, but requires a bit more work to prove formally:

Proposition 4 *If s^* is a strict dominant strategy equilibrium or if it uniquely survives IESDS, then s_i^* is a BR to $s_{-i}^*\,\forall\, i \in N$.*

proof: To see this, note that if s^* is a dominant strategy equilibrium, then it must uniquely survive IESDS, so it is enough to prove the proposition for strategies that uniquely survive IESDS. Suppose s^* uniquely survives IESDS, and choose some $i \in N$. Suppose in negation to the proposition that s_i^* is not a BR to s_{-i}^*. This implies that there exists some (maybe more than one) $s_i' \in S_i \diagdown \{s_i^*\}$ (this is the set S_i without the strategy s_i^*) such that $u_i(s_i', s_{-i}^*) > u_i(s_i^*, s_{-i}^*)$. Let $S_i' \subset S_i$ be the set of all such s_i'. Since s_i' was eliminated while s_{-i}^* was not (recall that s^* uniquely survives IESDS), there must be some s_i'' such that $u_i(s_i'', s_{-i}^*) > u_i(s_i', s_{-i}^*) > u_i(s_i^*, s_{-i}^*)$, implying that $s_i'' \in S_i'$. An induction argument on S_i' then implies that there exists a strategy $s_i''' \in S_i'$ that survives IESDS. But this is a contradiction to s^* being the unique survivor of IESDS. ■

Now that we have securely defined the notion of a best response, we need to think more seriously about the following question: to what should a player be playing a best response? Put differently, if my best response depends on what the other players are doing, then how should I choose between all the best responses I can possibly have? This is particularly pertinent since we are discussing static games, where players choose their actions without knowing what their opponents are choosing. To tackle this important question, we need to give players the ability to form *conjectures* about what others are doing.

7.2 Beliefs

Now suppose s_i' is a best response for player i against his opponents when they play s_{-i}', and assume for the moment that it is not a best response for any other profile of actions that i's opponents can choose. When would a rational player i choose to play s_i'? The answer must be clear: only when his *beliefs about other players' behavior* justifies the use of s_i', or in other words, when he believes that his opponents are going to play s_{-i}'.

Introducing the notion of beliefs, and actions that respond to beliefs, is central to the analysis of strategic behavior. When we have games in which players have dominant strategies, beliefs about the behavior of others had no role to play – all a player should care about is choosing his best strategy, and when he has a strictly dominant strategy his best strategy is *independent* of his opponents' play. But, when no such dominant strategies exist, a player must ask himself: "what do I think my opponents will do?" and the answer to this question should guide his own behavior.

Note, however, that when you try to guess the behavior of your opponents, you must take into account their guessing of your behavior, and any other player's behavior as well. To some extent, we employed such a thought process when we introduced the idea of IESDS: instead of asking what your opponents will do, you asked "what would a rational player *not do*?", and then, assuming that all players follow this process by common knowledge of rationality, we were able to make some prediction. In what follows, we will introduce another way of ruling out irrational behavior with a process that is, in many ways, the mirror of IESDS.

7.3 Rationalizability

Like IESDS, the concept of Rationalizability also uses the idea that the game and rational behavior are common knowledge. Instead, however, of asking "what would a rational player *not do*?", it asks, "what *might* a rational player *do*?" The answer is that a rational player will only select strategies that are a *best response to some profile of his opponents*. In other words, a strategy might be played by a rational player, if he can have *beliefs* that would justify the play of that strategy as a best response.

In turn, common knowledge of rationality implies that after employing this reasoning once, we can look at the resulting reduced game that includes only strategies that can be a best response, and then employ this reasoning again and again, in a similar way that we did for IESDS, in order to eliminate outcomes that should not be played by players who share a common knowledge of rationality.

The solution concept of rationalizability is defined precisely by iterating this thought process. The set of strategy profiles that survive this process are called the set of *rationalizable strategies*. (As with IESDS, we will not provide a formal definition since the introduction of mixed strategies is essential to do this.)

The Cournot Example Revisited

Consider the Cournot Duopoly example that demonstrated IESDS above with demand $p(q) = 100 - q$ and zero costs for both firms. Now that we have introduced the idea of a best response, it should be clear that in this example each firm's best response is immediately derived from the first order condition of its profit maximization problem, given by equation (6.1). In other words, if firm 1 believes that firm 2 will choose the quantity $\overline{q}_2$, then it should choose q_1 to satisfy:

$$q_1 = \begin{cases} \frac{100-\overline{q}_2}{2} & \text{if } 0 \leq \overline{q}_2 < 100 \\ 0 & \text{if } \overline{q}_2 \geq 100 \end{cases} .$$

Notice that this implies the following: firm 1 will *only* choose to produce quantities between 0 and 50, since there are no beliefs about $\overline{q}_2$ for which quantities above 50 are a best response. Thus, the first round of rationalizability implies that $q_1 \in [0, 50]$, and a symmetric argument works for firm 2. The next round imposes that $q_2 \in [0, 50]$, which means that the best reponse of firm 1 is to choose any quantity $q_1 \in [25, 50]$. Just as with IESDS, this process will continue on and on. The set of rationalizable strategies converges to a single quantity choice of $q_i = 33\frac{1}{3}$ for both firms.

A Discrete Example: the "p-beauty contest"

Consider the "p-beauty contest" as follows:

- Each players $i \in N = \{1, 2, ..., n\}$ choose an integer $s_i \in S_i = \{0, 1, ..., 20\}$
- The set of winners $W \subset N$ (may be one or more) are those players that are closest to 75% (here $p = 0.75$) of the average, and the rest are all losers. That is, the winners are

$$W = \left\{ \arg\min_{i \in N} \left| s_i - \frac{3}{4} \cdot \frac{1}{n} \cdot \sum_{i=1}^{N} s_i \right| \right\}$$

- Each player pays 1 to play the game, and winners split the pot equally among themselves. This implies that if there are $k \geq 1$ winners, each gets a payoff of $\frac{N-k}{k}$ (their share of the pot net of their own contribution) while losers get -1 (they lose their contribution).

Can we find strategies that can never be a best response? This is not a trivial task, but some simple insights can get us moving in the right direction. Since the objective is to guess a number closest to 75% of the average, this means that a player would want to guess a number that is generally (not always, as we will see) smaller than the highest numbers that other players may be guessing.

This implies that if there are strategies that are never a best response, they should be the higher numbers, and it is natural to start with 20 – can it be a best response? If you believe that the average is below 20, then 20 cannot be a best response – there will be a lower number that is. If the average is 20, that means that you and everyone else is choosing 20, and you would then split the pot with all the other players. If you believe this, and instead of 20 you choose 19, then you will win the whole pot for sure, regardless of the number of players, because for any number of players n 19 will be closer to $\frac{3}{4} \cdot \frac{1}{n} \cdot [(n-1) \cdot 20 + 19]$. This should convince you that 20 can never be a best response! Interestingly, 19 is not the *unique* best response to the belief that all others are playing 20, as the following exercise claims. The important point, however, is that 20 *cannot* be a best response to any beliefs a player can have.

Exercise 1 *Show that if player i believes that everyone else is choosing 20, any choice in the set $s_i \in \{11, 12, 13, 14, 15, 16, 17, 18, 19\}$ will be a best response, and lower numbers may also be a best response if n is not too large.*

Thus, only the numbers $S_i^1 = \{0, 1, ..., 19\}$ survive the first round of rationalizable behavior. Similarly, after each additional round we will "loose" the highest number until we go through 19 rounds and are left with $S_i^{19} = \{0, 1\}$, meaning that after 19 rounds of dropping strategies that cannot be a best response, we are left with two strategies that survive: 1 and 0. If $n > 2$, we cannot reduce this set further: if, for example, player i believe that all the other player's are choosing 1, then choosing 1 is a best response for him (as before, you need to do the calculation to fully convince yourself). Similarly, regardless of n, if he believes that everyone is choosing 0 then choosing 0 is his best response. Thus, we are able to predict using rationalizability, that players will not choose number greater than 1, and if there are only two players then we will predict that both will choose 0.

Will this indeed predict behavior? Only if our *assumptions* about behavior are correct!

In terms of existence, uniqueness and implications for Pareto optimality, rationalizability is practically the same as IESDS. It will sometimes have bite, and may even restrict behavior quite dramatically as in the examples above. But, as in the battle of the sexes, rationalizability will say "anything can happen" for many games.

Remark 4 *At this stage you may think that rationalizability and IESDS are not only two sides of the same coin, but down right identical. Indeed, in two player games they are if we allow for mixed strategies, as we will soon. Just to see why the concepts are not the same with pure strategies, consider the p-beauty contest above. We saw that choosing 20 is never a best response, so it would be eliminated in one round of Rationalizability. However, 20 is not strictly dominated by any other choice of an integer less than 20. To see this, note that 75% of the average will be some number between 0 and 15. This in turn implies that even though 20 can never be a winning number, it is also true that* any *number between 0 and 20* *can be a loser sometimes. Thus, for every number x between 0 and 20, there is a choice by the other players for which both 20 and x are losing numbers. But this in turn implies that 20 is not strictly dominated by any other integer between 0 and 20, so it cannot be eliminated by IESDS.*

Remark 5 *With more than two players IESDS and Rationalizability do not coincide. Showing this is a rather technical endeavor, and is beyond the scope of this text. See chapter 2 in Fudenberg and Tirole (1990) if you are interested.* L4

8

Pinning Down Beliefs: Nash Equilibrium

When we consider a game like the battle of the sexes, none of the concepts introduced above had any bite: dominant strategy equilibrium did not apply, and both IESDS and rationalizability could not restrict the set of reasonable behavior.

		Chris	
		O	F
Pat	O	2,1	0,0
	F	0,0	1,2

For example, we cannot rule out the possibility that Pat goes to the opera, while Chris goes to the football game since Pat may behave optimally to her belief that Chris is going to the opera, and Chris may behave optimally to his belief that Pat is going to the football game. But if we think of this pair of actions not only as actions, but as a system of actions and beliefs, then there is something of a dissonance: indeed the players are playing best responses to their beliefs, but their beliefs are wrong!

This is where we are about to make a huge leap in our requirements of a solution. For dominant strategy equilibrium, all we required is for people to be rational, but it applied very seldom. For IESDS and rationalizability, we demanded rationality, and common knowledge of it. Now, we will introduce a much more demanding concept, *Nash Equilibrium* introduced by John Nash – a Nobel Laureate, and the subject of a very successful Hollywood movie, *A Beautiful Mind* (based on the book by Sylvia Nasser).

In a nutshell, we can think of a Nash equilibrium is a system of beliefs and a profile of actions so that each player is playing a best response to his beliefs, and moreover, that players have correct beliefs. We will, for now, suppress any formal introduction of beliefs (which will play a much more important role later) and just focus on the strategies that players choose. In particular, a Nash equilibrium is a profile of strategies so that *each* player is choosing a best response to the strategies of *all other* players. Formally:

Definition 10 *The pure strategy profile, $s^* = (s_1^*, s_2^*, ..., s_n^*) \in S$ is a* ***Nash Equilibrium*** *if s_i^* is a BR to s_{-i}^*, for all $i \in N$,that is:*

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i', s_{-i}^*) \text{ for all } s_i' \in S_i \text{ and all } i \in N\,.$$

At the risk of being repetitive, lets emphasize what the requirements of a Nash equilibrium are:

1. Players are playing a *best response* to their beliefs
2. Players' beliefs about their opponents are *correct*

The first requirement is a direct consequence of rationality. It is the second requirement that is very demanding, and is a tremendous leap beyond the structures we have considered so far. It is one thing to ask people to behave rationally given their beliefs (play a best response), but a totally different thing to ask players to predict the behavior of their opponents correctly.

The again, there may be reasons to accept such a strong requirement if we allow for some reasoning that is beyond the physical structure of the game. For example, imagine that Pat is an influential gal – people try hard to make sure Pat is happy, and this is something that Pat knows well. In this case, Chris should believe, knowing that Pat is so influential, that she would expect him to go to the opera, and Pat's beliefs, knowing this, should be that Chris will try to please her and go to the opera. Indeed, (O, O) is a Nash equilibrium. However, notice that we can make the symmetric argument about Chris being an influential guy: (F, F) is also a Nash equilibrium. As the external game theorist, however, we should not say more than "one of these two outcomes is what we predict". (You should be able to convince yourself that no other pair of pure strategies is a Nash equilibrium.)

What about the other games we saw? In Prisoner's Dilemma, the unique Nash Equilibrium is (F, F). In the Cournot Duopoly game, the unique Nash Equilibrium is $(33\frac{1}{3}, 33\frac{1}{3})$, as we will see formally soon. Recall the following two-player discrete game we used to demonstrate IESDS:

	L	C	R
U	**4,3**	5,1	6,2
M	2,1	8,4	3,6
D	3,0	9,6	2,8

In it, the only pair of pure strategies that constitute a Nash equilibrium is (U, L), the same pair that survived IESDS.

The relationship between the outcomes we obtained earlier and the Nash equilibrium outcomes is no coincidence. There is a simple relationship between the concepts we previously developed and that of Nash equilibrium as the following proposition states clearly:

- **Proposition 5** *Consider a strategy profile* $s^* = (s_1^*, s_2^*, ..., s_n^*)$. *If* s^* *is either:*
  (1) a strict dominant strategy equilibrium;
  (2) the unique survivor of IESDS; or
  (3) the unique rationalizable strategy profile;
  then s^* *is the unique Nash Equilibrium.*

This proposition is simple to prove, and is left as an exercise. The intuition is of course quite straightforward: we know that if there is a strict dominant strategy equilibrium then it uniquely survives IESDS, and this in turn must mean that players are playing a best response to the other players' strategies.

As for our criteria to evaluate solution concepts, we can see from the Battle of the Sexes example that we may not have a unique Nash equilibrium. However, as the argument above alludes to, there is no reason to expect that we should – we may need to entertain other aspects, such as social norms, to make precise predictions. It turns out that for a rich set of games, as we will briefly discuss later, a Nash equilibrium always exists, which gives this solution concept its power – like IESDS and rationalizability, the solution concept of Nash is widely applicable. It will, however, usually lead to more refined predictions than those of IESDS and rationalizability.

From the prisoner's dilemma, we can easily see that Nash equilibrium does not guarantee Pareto optimality. Indeed, when people are left to their own devices, in some situations we need not expect them to make the best of it. This is where we can revisit the important restriction to *self enforcing outcomes*: our solution concepts took the game as given, and imposed rationality and common knowledge to try and see what players will choose to do. If they seek to maximize their own well being, they may hinder their ability to achieve socially optimal outcomes.

8.1 Applications: Duopoly

8.1.1 Cournot Duopoly

Let's revisit the Cournot game with demand $P = 100 - q$ and cost functions $c_i(q_i) = c_i \cdot q_i$ for firms $i \in \{1, 2\}$. The maximization problem that firm i faces when it believes that its opponent chooses quantity q_j is,

$$\max_{q_i} \; \pi(q_i, q_j) = (100 - q_i - q_j) \cdot q_i - c_i \cdot q_i \; .$$

Recall that the best response for each firm is given by the first-order condition,

$$BR_i(q_j) = \frac{100 - q_j - c_i}{2}.$$

This means that each firm chooses quantities as follows:

$$q_1 = \frac{100 - q_2 - c_1}{2} \quad \text{and} \quad q_2 = \frac{100 - q_1 - c_2}{2}\,. \tag{8.1}$$

When do we have a Nash equilibrium? Precisely when we find a pair of quantities, (q_1, q_2) that are *mutual best responses*. This occurs exactly when we solve both best response functions (8.1) simultaneously. The following diagram shows the solution for $c_1 = c_2 = 0$, in which case the unique Nash equilibrium is $q_1 = q_2 = 33\frac{1}{3}$.

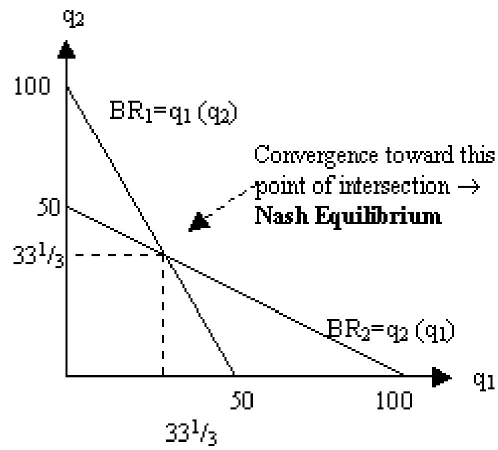

Exercise 2 *Suppose there are n firms in the Cournot oligopoly model. Let q_i denote the quantity produced by firm i, and let $Q = q_i + \cdots + q_n$ denote the aggregate production. Let $P(Q)$ denote the market clearing price (when demand equals Q) and assume that inverse demand function is given by $P(Q) = a - Q$ (where $Q < a$). Assume that firms have no fixed cost, and the cost of producing quantity q_i is $c \cdot q_i$ (all firms have the same marginal cost, and assume that $c < a$). (i) Model this as a Normal form game; (ii) What is the Nash (Cournot) Equilibrium of the game where firms choose their quantities simultaneously? (iii) What happens to the equilibrium price as n approaches infinity?*

8.1.2 Bertrand Duopoly

The Cournot model assumed that the firms choose quantities, and the market price adjusts to clear the demand. However, one can argue that firms often set prices, and let consumers choose where to purchase from, rather than setting quantities and waiting for the market price to equilibrate demand. We now consider the game where each firm posts a price for their, otherwise identical goods. This was the situation modelled and analyzed by Joseph Bertrand in 1883.

As before, assume that demand is given by $p = 100 - q$ and cost functions $c_i(q_i) = 0$ for firms $i \in \{1, 2\}$ (zero costs). Clearly, we would expect buyers to all buy from the firm whose price is the lowest. What happens if there is a tie? Let's assume that the market splits equally between the two firms. This gives us the following normal for of the game:

- $N = \{1, 2\}$
- Strategy sets are $S_i = [0, \infty]$ for $i \in \{1, 2\}$ and firms choose prices $p_i \in S_i$
- To calculate payoffs, we need to know what the quantities will be. The quantities are given by:

$$q_i(p_i, p_j) = \begin{cases} 100 - p_i \text{ if } p_i < p_j \\ 0 \text{ if } p_i > p_j \\ \frac{100-p_i}{2} \text{ if } p_i = p_j \end{cases}$$

which in turn means that the payoff function is given by:

$$u_i(p_i, p_j) = \begin{cases} (100 - p_i) \cdot p_i & \text{if } p_i < p_j \\ 0 & \text{if } p_i > p_j \\ \frac{100-p_i}{2} \cdot p_i & \text{if } p_i = p_j \end{cases}$$

We now need to calculate the best response functions of both firms. To do this, we will first start with a slight modification, that is motivated by reality: assume that prices cannot be any real number, but are limited to be increments of some small number, say $\varepsilon > 0$. That is, prices are assumed to be in the set $\{0, \varepsilon, 2\varepsilon, 3\varepsilon ...\}$. For example, $\varepsilon = 0.01$ if we are considering cents as the price increment,[1] and the strategy set will be $\{0, 0.01, 0.02, ...\}$. We will soon introduce smaller denominations, and will look at what happens when this increments become infinitely small and approach zero.

We can now explore the best response of a firm by exhausting the relevant situations that it can face. Assume first that p_j is very high, above 50. Then, firm i can set the monopoly (profit maximizing) price of 50 and not face any competition,

[1] Notice, for example, that in gas stations gallons are often quoted in prices that include one-tenth of a cent.

which is clearly what i would choose to do. Now assume that $50 > p_j > 0.01$. Firm i can choose one of three options: either set $p_i > p_j$ and get nothing, set $p_i = p_j$ and split the market, or set $p_i < p_j$ and get the whole market. It is easy to see that of these three, firm i wants to just undercut firm j and capture the whole market, thus setting a price of $p_i = p_j - 0.01$.[2] When $p_j = 0.01$ then these three options are still there, but undercutting means setting $p_i = 0$, which is the same as setting $p_i > p_j$ and getting nothing. Thus, the best reply is setting $p_i = p_j = 0.01$ and splitting the market. Finally, if $p_j = 0$ then any choice of price will give firm i zero profits, and therefore anything is a best response. In summary:

$$BR_i(p_j) = \begin{cases} 50 & \text{if } p_j > 50 \\ p_j - 0.01 & \text{if } 50 \geq p_j > 0.01 \\ 0.01 & \text{if } p_j = 0.01 \\ p_i \in \{0, 0.01, 0.02, 0.03, ...\} & \text{if } p_j = 0 \end{cases}$$

Now given that firm j's best response is exactly symmetric, it should not be hard to see that there are two Nash Equilibria that follow immediately from the form of the best response functions: The best response to 0.01 is 0.01, and a best response to 0 is 0. Thus, the two Nash equilibria are,

$$(p_1, p_2) \in \{(0, 0), (0.01, 0.01)\}.$$

It is worth pausing here for a moment to prevent a rather common point of confusion, which arises often when a player has more than one best response to a certain action of his opponents. In this example, when $p_2 = 0$, player 1 is indifferent between any price he chooses: if he splits the market with $p_1 = 0$ he gets half the market with no profits, and if he sets $p_1 > 0$ he gets no customers and has no profits. One may be tempted to jump to the following conclusion: if player 2 is choosing $p_2 = 0$, then any choice of p_1 together with player 2's zero price will be a Nash equilibrium. This is incorrect. It is true that player 1 is playing a best response with any choices, but if $p_1 > 0$ then $p_2 = 0$ is *not* a best response as we

[2] To see this, if we have some $p_j > 0.01$ then by setting $p_i = p_j$ firm i gets $\pi_i = \frac{100p_j - p_j^2}{2}$ while if it sets $p_i = p_j - 0.01$ it will get $\pi_i' = 100(p_j - 0.01) - (p_j - 0.01)^2$. If we calculate the difference between the two we get that $\pi_i' - \pi_i = 50.02p_j - \frac{1}{2}p_j^2 - 1.000\,1$, which is positive at $p_j = 0.02$, and this difference has a positive derivative for any $p_j \in [0.02, 50]$.

can observe from the analysis above. Thus, having one firm choose a price of zero while the other is not cannot be a Nash equilibrium.

Comparing the Bertrand game outcome to the Cournot game outcome is an interesting exercise. Notice that when firms choose quantities (Cournot), the unique Nash equilibrium when costs were zero had $q_1 = q_2 = 33\frac{1}{3}$. A quick calculation shows that $p = 33\frac{1}{3}$ and each firm makes a profit of $\$1,111.11$. When instead these firms compete on prices, the two possible equilibria have either zero profits when both choose zero prices, or negligible profits (about 50 cents) when they each choose a price of \$0.01. Interestingly, for both the Cournot and Bertrand games, if we only had one player, he would choose the monopoly price of \$50 and earn a profit of \$2,500.[3]

The message of this analysis is quite striking: one firm may have monopoly power, but when we let one more firm compete, and they compete with prices, then the market will behave competitively – if both choose a price of zero, price will equal marginal costs! Notice that if we add a third and fourth firm, this will not change the outcome; prices will have to be zero (or practically zero at 0.01) for all firms in the unique Nash (Bertrand) equilibrium. This is not the case for Cournot competition (exercise above).

A quick observation should lead you to realize that if we let ε be smaller than one cent, the conclusions above will be sustained, and we will have two Nash equilibria. One with $p_1 = p_2 = 0$, and one with $p_1 = p_2 = \varepsilon$. It turns out that these two equilibria not only become closer in profits as ε gets smaller, but we we take ε to zero and assume that prices can be chosen as any real number, we get a very "clean" result: the unique Nash equilibrium will have prices equal to marginal costs, implying a competitive outcome.

- **Proposition 6** *For $\varepsilon = 0$ (prices can be any real number) there is a unique Nash equilibrium: $p_1 = p_2 = 0$.*

proof: First note that we can't have a negative price in equilibrium – a firm offering it will lose money (pay the consumers to take its goods!). We need to show that we can't have a positive price in equilibrium. We can see this

[3] This is achieved by maximizing profits which are $\pi = (100 - q)q$.

in two steps:
(i) If $p_1 = p_2 = \hat{p} > 0$, each would benefit from changing $\hat{p} - \varepsilon$ (ε very small) and get the whole market for almost the same price.
(ii) If $p_1 > p_2 \geq 0$, p_2 would want to deviate to $p_1 - \varepsilon$ (ε very small) and earn higher profits.
It is easy to see that $p_1 = p_2 = 0$ is an equilibrium: both are playing a best response. ■

Exercise 3 *If the cost function was $c \cdot q_i$ for each firm then the unique Nash equilibrium would be $p_1 = p_2 = c$. Convince yourself of this as an easy exercise.*

We will now see an interesting variation of the Bertrand game. Assume that $c_i(q_i) = c_i \cdot q_i$ represents cost of the firm as before. Now, however, let $c_1 = 1$ and $c_2 = 2$ so that the two firms are not identical: form 1 has a cost advantage. Let the demand still be $p = 100 - q$.

Now consider the case with discrete price jumps with $\varepsilon = 0.01$. As before, there are still two Nash Equilibria:

$$(p_1, p_2) \in \{(1.99, 2.00), (2.00, 2.01)\}.$$

or more generally, $(p_1, p_2) \in \{(2 - \varepsilon, 2), (2, 2 + \varepsilon)\}$.

Exercise 4 *Show that only these are the Nash equilibria of this game.*

Now we can ask ourselves, what happens if $\varepsilon = 0$? If we would think of using a "limit" approach to answer this, then we may expect a similar result to the one we saw before, namely, that we get one equilibrium which is the limit of both, and in this case it would be $p_1 = p_2 = 2$.

But is this an equilibrium? Interestingly, the answer is no! To see this, consider the best response of firm 1. Its payoff function is not continuous when firm 2 offers a price of 2. The profit function of firm 1, as a function of p_1 when $p_2 = 2$, is depicted in the following figure:

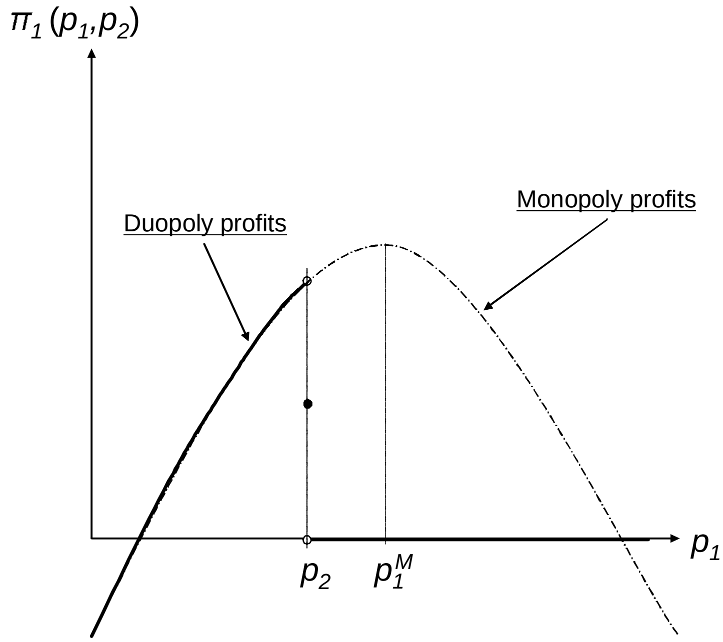

The figures first draws the profits of firm 1 if it were a monopolist with no competition, and if this were the case it would charge its monopoly price $p_1^M = 55.5$.[4] If firm 2 charged more than the monopoly price, this would have no impact on the choice of firm 1 – it will still charge the monopoly price. If, however, firm 2 charges a price p_2 that is less than the monopoly price then there is a discontinuity in the profit function of firm 1: as its price p_1 approaches p_2 it's profits rise, but when it hits p_2 exactly it will split the market and experience its profits dropping by half.

This discontinuity causes firm 1 to *not have* a best response when $p_2 < 55.5$. Firm 1 wants to set a price as close to p_2 as it can, but does not want to reach p_2 because then it splits the market and gets a "jump" down in profits. Once its price goes above 2 then firm 1's profits drop further to zero.

Indeed, this is an example where a Nash equilibrium does not exist. The reason is precisely the fact that firm 1 does not have a "well behaved" payoff function, which in turn causes it to not have a well defined best response function, a consequence

---

[4] The maximization here is for the profit function $\pi_1 = (100 - p)(p - c_1)$ where $c_1 = 1$.

of the discontinuity in the profit function. We will discuss this a bit more later, but not in too much depth.

8.2 Pure-Strategy NE in a Matrix: A Simple Method

In this short section, we go over a trivial, yet fool proof way to find Nash equilibria in matrix games when a pure strategy Nash equilibrium exists. Consider the following two person finite game in matrix form:

$$\begin{array}{c|c|c|c|}
 & L & C & R \\ \hline
U & 7,7 & 4,2 & \overline{1,8} \\ \hline
M & 2,4 & \underline{\overline{\mathbf{5,5}}} & \underline{2,3} \\ \hline
D & \underline{8,1} & \overline{3,2} & 0,0 \\ \hline
\end{array}$$

It is easy to see that no strategy is dominated, and thus we cannot say anything with IESDS or rationalizability. However, a pure strategy Nash equilibrium exists. To find it, we use a simple method that captures the fact that any Nash equilibrium must involve each player playing a best response to his opponent.

Step 1: For every *column*, under-line the pair of payoffs for which player 1 is playing a best-response

What step 1 does is to identify the best response of player 1 for each of the pure strategies (columns) of player 2. For instance, if player 2 is playing L, then player 1's best response is D, and we underline the payoffs associated with this row in column 1. In other words, for the choices (D, L) player 1 is playing a best response. This is also true for $(M, , C)$ and for (M, R).

Step 2: For every *row*, over-line the pair of payoffs for which player 2 is playing a best-response

Step 1 similarly identifies the best response of player 2 for each of the pure strategies (rows) of player 1. For instance, if player 1 is playing D, then player 2's best response is C, and we over-line the payoffs associated with this column in row 3. In other words, for the choices (D, C) player 2 is playing a best response. This is also true for $(M, , C)$ and for (U, R).

Step 3: If an entry has both an under- and over-line, it is the outcome of a Nash Equilibrium in pure strategies.

In this example, we find that (M, C) is the unique pure strategy Nash equilibrium – it is the only pair of pure strategies for which both players are playing a best response. If you apply this to the battle of the sexes, you will find both pure strategy Nash equilibria, (O, O) and (F, F). For the prisoner's dilemma only (F, F) will be identified.

As mentioned above, this trivial method will only find pure strategy Nash equilibria, that is, pairs of actions where each player is choosing *a particular one* of his possible actions. Now consider the following classic game called "Matching Pennies". Players 1 and 2 both put a penny on a table simultaneously. If the two pennies come up the same (heads or tails) then player 1 gets both, otherwise player 2 does. We can represent this in the following matrix:

		Player 2	
		H	T
Player 1	H	$\underline{1}, -1$	$-1, \overline{1}$
	T	$-1, \overline{1}$	$\underline{1}, -1$

Clearly, the method introduced above does not find a pure strategy Nash equilibrium – given a belief that player 1 has, he always wants to match it, and given a belief that player 2 has, he would like to choose the opposite orientation for his penny. Does this mean that a Nash equilibrium fails to exist? Not if we consider a richer set of possible behaviors, as we will now see. L5