Chapter 7, 8, and 9 Notes

These notes essentially correspond to parts of chapters 7, 8, and 9 of Mas-Colell, Whinston, and Green. We are not covering Bayes-Nash Equilibria. Essentially, the Economics Nobel Prize was given in 1994 to John Nash (for Nash Equilibrium), Reinhard Selten (for subgame perfect NE), and John Harsanyi (for Bayes-Nash Equilibrium). We cover two of the three, and you get Bayes-Nash Equilibria next semester.

1 Introduction

Up to now in this course we have been concerned with some type of optimization. Either a maximization problem (UMP and PMP) or a minimization problem (EMP and CMP), or an economy where consumers and producers solve optimization problems. Further, the decision made by a particular consumer did not affect the outcomes – if the economy was in equilibrium, and the prices set, then each individual simply made their choices and received their utility from those choices. There was no consideration of "what the other person was doing" and how this might affect my outcome or payoff. While we will still have consumers and producers optimizing, now we examine decisions where the outcomes that occur are a function of both the individual's decision and some other individual's (or multiple individuals) decision.

To solve problems of this type we will use game theory or the theory of games. "Game Theory" is kind of an oxymoron, like "Jumbo Shrimp", or, if you are John Kerry trying to make a poor joke, like "Military Intelligence". The word "game" evokes images of fun while the word "theory" evokes images of something a little more abstract or difficult. The most prominent early game theorists were John von Neumann and Oskar Morgenstern (we have already talked about vNM utility functions). von Neumann was the driving force behind the mathematics, but Morgenstern was instrumental in getting the book, Theory of Games and Economic Behavior, published in the mid-40s (On the military intelligence thing, if I recall correctly von Neumann and many other prominent academics were working for RAND which was working for the military during WWII. Let's put it this way – if you were doing a fantasy draft of intellectual minds, like one would do a fantasy draft for baseball or football, von Neumann would have been a first round pick in the 1940s). John Nash is probably the most famous game theorist (that's what happens when you are portrayed by Russell Crowe in a movie), and we will discuss his solution concept at length. But game theory provides structure for solving games where there are interdependencies among the participants in the game. It can be used to analyze actual games (Chess – which, incidentally, if you can solve Chess you will be famous, I assure you; Baseball; Candy Land; whatever) as well as things you may not think are games (such as an oligopoly market or committee voting). Eventually we will study oligopoly markets, but for now we will just discuss the basics of games. Each game consists of 4 components. We can use Chess as an example:

1. Players – Who actually plays the game? There are two players in Chess, one who controls the White pieces and one who controls the Black pieces. Note that players refer to those people who actually make decisions in the game.

2. Rules – Who makes what decisions or moves? When do they make the moves? What are they allowed to do at each move? What information do they know? In a standard Chess game, White moves first and there are 20 moves that White can make (8 pawns that can move either one or two spaces ahead, and 2 knights that can move to one of 2 different spots on the board). Players alternate turns, so that Black also has 20 moves that can be made on his first turn. Furthermore, there are restrictions on how the pieces can move, how pieces are removed and returned to the board, how a winner is determined, how long a player has to make a move – in short, there are a lot of rules to Chess.

3. Outcomes – What occurs as a result of the rules and the decisions players make? At the end of a Chess match one of three things occurs – White wins, Black wins, or there is a draw. Those are the end results of the game. Much simpler than the rules.

4. Payoffs – What utility is assigned to each of the outcomes? Essentially each player has a utility function over outcomes and acts in a manner to best maximize utility, taking into consideration that the other player is doing the same. It does not have to be the case that "winning" has a higher utility than "losing". It may be that one's payoff is tied to who the other players are. If the Chess match is a professional or amateur match and you can win money (or fame) by winning the match, then typically winning will have a higher payoff than losing. However, if you are playing a game with your child or sibling and you are attempting to build their self-esteem then perhaps losing has a higher payoff. Basically, there is a utility function that is a function of all the relevant variables and this utility function determines the players payoffs. In most cases we will simply assume the payoffs are interchangeable with the outcomes, so that specifying a payoff specifies an outcome.

If there is only one player then it is not a game but a decision. Decisions are easy to solve – simply make a list of available actions to the player and then choose the action that gives the player the highest payoff. This is like our consumer maximization problem, although there are many, many decisions a consumer could make in an economy with L goods. But the consumer lists all those combinations available to him or her (the budget constraint acts as a rule or restriction on what is available) and then makes a choice about which available bundle maximizes utility. Note that there can be one-player games if there is some uncertainty involved. Take Solitaire as an example. There is only one player making an active decision, but there is a second "player", which we would call "nature" or "random chance". The player makes an active decision to make a particular move, and then nature makes a move regarding the next card to be shown. But we are getting ahead of ourselves.

Games are slightly more complicated than decisions because the other player's decisions must be taken into consideration as they affect the outcomes and payoffs to all players. We will begin by considering simultaneous move games and then move to a discussion of sequential games. For now, we consider games with no uncertainty over payoffs or randomness due to nature and we assume common knowledge. Common knowledge means that player 1 knows what player 2 knows, and player 2 knows that player 1 knows what player 2 knows, and the player 1 knows that player 2 knows that player 1 knows what player 2 knows, ad infinitum. Now, a few formalities:

Definition 1 *Let* $\not{H}_i$ *denote the collection of player i's information sets,* $\tilde{A}$ *the set of possible actions in the game, and* $C(H) \subset \tilde{A}$ *the set of actions possible at information set* H. *A strategy for player* i *is a function* $s_i : \not{H}_i \to \tilde{A}$ *such that* $s_i(H) \in C(H)$ *for all* $H \in \not{H}_i$.

First of all, what is an information set? An information set is what a player knows about the *moves* that the other player has made in the game. Thus, if there are 2 players and they make moves simultaneously then player i knows nothing about the move made by player j and player i's information is only the structure of the game. If the players are playing a sequential game such as Chess, then when the Black player makes his first move he knows which of the 20 moves the White player made. So, his information set is that White moved piece X to square Y, and he can now disregard the other 19 moves that White could have made initially. He still knows that White could have made these moves, and maybe that tells Black something about the strategy White is using, but the simple fact is that White made a move, Black saw it, and now Black must make a move based on a Chess board that looks a particular way after White makes his move.

Now, about strategies. In a very simple game, which we will get to shortly, it may be that both players only make one move, the players move simultaneously, and then the game ends. In this case a strategy is just one "decision" or move for a player. In that case, the decision made specifies what the player will do in every possible contingency that might arise in the game. However, consider Chess. A strategy for Chess is much more complex. Consider the Black player's first move. White can make 20 different opening moves. Black must provide an action for each of these potential opening moves. Thus, there are 20 actions that must be specified by Black, and that is just for his first move!!! After White and Black both make their initial moves, White now has to specify 400 actions for his second move (20 potential opening moves by White times 20 potential opening moves by Black). And now we are only at the third move of the entire game. This is why Chess has not yet been solved. Thus, a strategy for a player is a complete contingent plan for that player.

2 Simultaneous move games

We will begin by considering games in which players move at the same time. These games could be truly simultaneous, or it could be that the players make actions at different times but that neither player knows of the actions taken by the other.

2.1 Pure strategies

Consider the following story:

You and someone else in the class have been charged with petty theft. You are strictly interested in your own well-being, and you prefer less jail time to more. The two of you are isolated in different holding cells where you will be questioned by the DA. The DA comes in and makes you an offer. The DA says that if you confess and your partner confesses that you will both be sentenced to 8 months in jail. However, if neither one of you confess, then both of you will do the minimum amount of time and be out in 2 months. But if you confess and your partner does not confess then you can go home free (spend zero months in jail) and your partner gets to spend 12 months in jail (where 12 months is the maximum amount of time for this crime). The DA also tells you that your partner is offered the same plea bargain and that if he/she confesses and you don't, then you will be sentenced to 12 months in jail and your partner will go home free. Do you confess or not confess?

This is the classic prisoner's dilemma story. We will use a matrix representation of the game (it is a game – there are players, actions, outcomes, and payoffs) to analyze it. Matrix representation of the game is also known as the normal form or strategic form of the game.

Definition 2 *For a game with I players, the normal form representation Γ_N specifies for each player i a set of strategies S_i (with $s_i \in S_i$) and a payoff function $u_i(s_1, ..., s_I)$ giving the vNM utility levels associated with the (possibly random) outcome arising from strategies $(s_1, ..., s_I)$. Formally, we write $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$.*

Basically, S_i is the set of all strategies available to player i, and s_i is a particular strategy chosen from S_i by player i. The payoff function $u_i(\cdot)$ is a function of the strategies chosen by all I players. This is why the normal form representation is also known as the strategic form – there is no mention of the order of moves, just a list of strategies that each player can take. When we discuss sequential games we will discuss how any sequential game can be represented in strategic or normal form. For now, consider the matrix in Figure 1 for the prisoner's dilemma game described above.

In a two-player game we have one player who is labeled the "row player" and another player who is labeled the "column player". In this case, Prisoner 1 is the row player and Prisoner 2 is the column player. The row player's strategies are listed along the rows, while the column player's strategies are listed across the columns. Each player has 2 strategies, confess or don't confess. In the cells of the matrix we put the payoffs from the choice of these strategies – by convention, the row player's payoff is listed as the payoff on the left and the column player's payoff is the one on the right. If both prisoners confess then they each spend 8 months in prison. If both players do not confess then they each spend 2 months in prison. If Prisoner 1 confesses and Prisoner 2 does not, then Prisoner 1 spends 0 months in prison and Prisoner 2 spends 12 months in prison. The opposite is true if Prisoner 1 does not confess and Prisoner 2 does confess. The matrix form lists all the strategies available to each player and the payoffs associated with the player's choice of strategies. Formally, $S_i = \{Confess, Don't\ Confess\}$ for $i = 1, 2$, with s_1 and s_2 either Confess or Don't Confess (the specific strategy, not the set of strategies). The payoff $u_1(confess, don't\ confess) = 0$, the payoff $u_1(confess, confess) = -8$, the payoff $u_1(don't\ confess, don't\ confess) = -2$, and the payoff $u_1(don't\ confess, confess) = -12$. All of the elements of a normal form game are represented in the matrix.

Now, how do we solve the game? We are looking for a Nash Equilibrium (NE) of the game. A Nash Equilibrium of the game is a set of strategies such that no player can unilaterally deviate from his chosen strategy and obtain a higher payoff. There are a few things to note here. First, a Nash Equilibrium is a set

		Prisoner 2	
		Confess	Don’t Confess
Prisoner 1	Confess	-8 , -8	0 , -12
	Don’t Confess	-12 , 0	-2 , -2

Figure 1: Matrix representation for the prisoner’s dilemma.

of STRATEGIES, and not payoffs. Thus, a NE of the game may be $\{Confess, Don't\ Confess\}$ or it may be $\{Confess, Confess\}$ or you might write that Prisoner 1 chooses Don't Confess and Prisoner 2 chooses Don't Confess if you are not into the whole brevity thing. I am not particular about notation for these simple games, but if you write down the NE is something like $\{-2, -2\}$, and $\{-2, -2\}$ represents a payoff and not a strategy, it is very likely that I will not look at the rest of the answer. Now, if the strategy is actually a number, then it is fine to write down a number, such as Firm 1 chooses a quantity of 35 and Firm 2 chooses a quantity of 26, but please remember that NE are STRATEGIES, not payoffs. Second, when we consider NE we look at whether or not one player can unilaterally deviate from the chosen strategies of all the players to increase his payoff. It is possible that multiple players would deviate and this would increase their payoffs, but we are going to hold the chosen strategies of the other players constant and see if a particular player would deviate. Formally, we define a Nash Equilibrium as:

Definition 3 *A strategy profile $s = (s_1, ..., s_I)$ constitutes a Nash Equilibrium of game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if for every $i = 1, ..., I$,*

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$$

for all $s_i' \in S_i$.

Now, how to solve the simple game of the Prisoner's Dilemma. If Prisoner 1 was to choose Confess, and Prisoner 2 knew this, what would Prisoner 2 choose? Prisoner 2 would choose Confess. If Prisoner 1 was to choose Don't Confess, and Prisoner 2 knew this, what would Prisoner 2 choose? Prisoner 2 would still choose Confess. We can show the same result for Prisoner 1 holding Prisoner 2's choice of strategy constant. Thus, the NE of the Prisoner's Dilemma is Prisoner 1 chooses Confess and Prisoner 2 chooses Confess. When there is a simple matrix, it is easy enough to circle the payoffs as we did in class.

There are a few things to note here. One is that both players choice of strategy does not depend on what the other does. Regardless of what Prisoner 1 does Prisoner 2 should choose Confess, and the same is true for Prisoner 1. Thus, both players have a strictly dominant strategy in this game.

Definition 4 *A strategy $s_i \in S_i$ is a strictly dominant strategy for player i in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if for all $s_i' \neq s_i$ we have:*

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

The strategy is weakly dominant if $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$ (note that the difference is that the inequality is not strictly greater than, but greater than or equal to).

Thus, one of the first things to look for is a strictly dominant strategy for any players. If a player has a strictly dominant strategy, then that simplifies the solution of the game tremendously, because all of the other players *SHOULD* know that the player with the strictly dominant strategy will not choose anything other than that strategy.[1]

A related concept is that of a strictly *dominated* strategy (note the difference between *dominant* and *dominated*).

Definition 5 *A strategy $s_i \in S_i$ is strictly dominated for player i in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if there exists another strategy $s_i' \in S_i$ such that for all $s_{-i} \in S_{-i}$*

$$u_i(s_i', s_{-i}) > u_i(s_i, s_{-i}).$$

The strategy s_i is weakly dominated if $u_i(s_i', s_{-i}) \geq u_i(s_i, s_{-i})$.

Thus, a strictly dominated strategy is one that a player would not choose regardless of the strategies chosen by the other players. In the Prisoner's Dilemma game, the strategy "Do not confess" is strictly dominated by the strategy "Confess".

[1] Whether people actually realize this or act in a manner that suggests they realize this is debateable.

2.2 Incorporating Mixed Strategies

It is possible that a player chooses not to play a pure strategy from the set $\{S_i\}$, but to *randomize* over available strategies in $\{S_i\}$. Thus, a player may assign a probability to each strategy, adhering to the common laws of probability (all probabilities sum to 1, no probabilities greater than 1 or less than 0). Note that the idea is to randomize, which we will discuss a little more in depth momentarily.

Definition 6 *Given player $i's$ (finite) pure strategy set S_i, a mixed strategy for player i, $\sigma_i : S_i \rightarrow [0,1]$, assigns to each pure strategy $s_i \in S_i$ a probability $\sigma_i(s_i) \geq 0$ that it will be played, where $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$.*

Suppose player i has M pure strategies in set $S_i = \{s_{1i}, ..., s_{Mi}\}$. Player $i's$ set of possible mixed strategies can therefore be associated with the points of the following simplex:

$$\begin{aligned} \Delta(S_i) &= \{\sigma_{1i}, ..., \sigma_{Mi}\} \in \mathbb{R}^M : \sigma_{mi} \geq 0 \\ \text{for all } m &= 1, ..., M \text{ and } \textstyle\sum_{m=1}^{M} \sigma_{mi} = 1 \end{aligned}$$

Hence, $\Delta(S_i)$ is simply a mixed extension of S_i, where pure strategies are the degenerate probability distributions where $\sigma_{ji} = 1$ for some strategy j.

When players randomize over strategies the induced outcome is random. In the normal form game the payoff function for i is $u_i(s)$. This payoff function is a vNM type, so that player $i's$ expected utility from a mixed strategy profile $\sigma = (\sigma_1, ..., \sigma_I)$ is $E_\sigma[u_i(s)]$, with the expectation taken with respect to the probabilities induced by σ on pure strategy profiles $s = (s_1, ..., s_I)$. Denote the normal form representation as $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$, which includes mixed and pure strategies.

Now that we have discussed the concept of mixed strategies, let us formalize the concepts of best response and Nash Equilibrium.

Definition 7 *In game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$, strategy σ_i is a best response for player i to his rivals' strategies σ_{-i} if*

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i}) \text{ for all } \sigma_i' \in \Delta(S_i)$$

Strategy σ_i is never a best response if there is no σ_{-i} for which σ_i is a best response.

Essentially, a mixed strategy σ_i is a best response to some choice of mixed strategies for the other players σ_{-i} if the utility from σ_i is at least as large as the utility from any other available mixed strategy σ_i'. We can also think about best responses in terms of a best response correspondence (this will prove useful when studying the Cournot model and discussing existence of pure strategy Nash Equilibria). We will focus on pure strategies here, rather than mixed, for reasons which will be made clear later.

Definition 8 *A player's best response correspondence $b_i : S_{-i} \rightarrow S_i$ in the game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$, is the correspondence that assigns to each $s_{-i} \in S_{-i}$ the set*

$$b_i(s_{-i}) = \{s_i \in S_i : u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \text{ for all } s_i' \in S_i\}.$$

Thus, a player's best response correspondence will tell us which strategy (strategies) do best against the other player's strategies. Now, a formal definition of Nash Equilibrium:

Definition 9 *(Nash Equilibrium allowing mixed strategies) A mixed strategy profile $\sigma = (\sigma_1, ..., \sigma_I)$ constitutes a Nash Equilibrium of game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ if for every $i = 1, ..., I$,*

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i}) \text{ for all } \sigma_i' \in \Delta(S_i)$$

All this says is that all I players are playing a best response to each other. Note that this encompasses pure strategies since they are simply degenerate mixed strategies. However, pure strategies tend to be more interesting than mixed strategies, so we will restate the definition in terms of pure strategies. We will also use the concept of a best response function.

Definition 10 *(Nash Equilibrium in pure strategies) A strategy profile $(s_1, ..., s_I)$ is a Nash Equilibrium of game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if and only if $s_i \in b_i(s_{-i})$ for $i = 1, ..., I$.*

Thus, a set of pure strategies is a Nash Equilibrium if and only if the strategies are best responses to one another.

Let's talk about a particular game, Matching Pennies. There are two players who move simultaneously in this game. Each player places a penny on the table. If the pennies match (both heads or both tails) then Player 1 receives a payoff of 1 and Player 2 receives a payoff of (-1). If the pennies do not match (one heads and one tails), then Player 1 receives a payoff of (-1) and player 2 receives a payoff of 1. The matrix representation of the game is here:

		Player 2	
		Heads	Tails
Player 1	Heads	$1, -1$	$-1, 1$
	Tails	$-1, 1$	$1, -1$

Note that there is no pure strategy Nash Equilibrium to this game. However, there may be a mixed strategy Nash Equilibrium to the game. For now, suppose that Player 1 chooses Heads 50% of the time and Tails 50% of the time. Player 2's expected payoff from ANY strategy (mixed OR pure) is 0. If Player 2 chooses Heads with probability 1, then Player 2's payoff is $1 * 50\% + (-1) * 50\% = 0$. It is the same if Player 2 chooses Tails with probability 1, or if Player 2 chooses a 50/50 mix, or a 75/25 mix, or a 25/75 mix. Thus, Player 1's choice of Heads 50% of the time and Tails 50% of the time has made Player 2 indifferent over any of his strategies. Now, is Player 1 choosing Heads 50% of the time and Tails 50% of the time and Player 2 choosing Tails 100% of the time a Nash Equilibrium of this game? No, because if Player 2 were to choose Tails 100% of the time then Player 1 would wish to choose Tails 100% of the time (or at least shift the probabilities so that choosing Tails is weighted more heavily than choosing Heads). Thus, for a set of mixed strategies to be a Nash Equilibrium BOTH (or all) players must be making each other indifferent to all strategies (almost – all pure strategies that the player includes in the mixing distribution). Even if Player 2 chose Tails 51% of the time and Heads 49% of the time Player 1 could still do better by choosing Tails 100% of the time. These best response functions are shown in Figure 2. This idea is formalized below:

Proposition 11 *Let $S_i^+ \subset S_i$ denote the set of pure strategies that player i plays with positive probability in mixed strategy profile $\sigma = (\sigma_1, ..., \sigma_I)$. Strategy profile σ is a Nash Equilibrium in game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ if and only if for all $i = 1, ..., I$*

1. *$u_i(s_i, \sigma_{-i}) = u_i(s_i', \sigma_{-i})$ for all $s_i, s_i' \in S_i^+$*
2. *$u_i(s_i, \sigma_{-i}) \geq u_i(s_i', \sigma_{-i})$ for all $s_i \in S_i^+$ and all $s_i' \notin S_i^+$*

This proposition states that all players must be indifferent over all their pure strategies over which they assign positive probability in their mixed strategy and that the utility from those strategies not included in their mixed strategy must be equal to or less than the utility from a pure strategy included in their mixed strategy. Let $S_i = \{A, B, C\}$ and $S_i^+ = \{A, B\}$. Then the utilities from the pure strategies A and B when played against the mixed strategy σ_{-i} must be equal, but the utility from pure strategy C may be less than or equal to that of pure strategy A or B. Because of this proposition, we need only check indifference among pure strategies and not all possible mixed strategies. If no player can improve by switching from the mixed strategy σ_i to a pure strategy s_i then the strategy profile σ is a mixed strategy Nash Equilibrium.

How to actually go about finding these mixing probabilities when there is a small number of pure strategies. Consider the Matching Pennies game. Let σ_{1H} be the probability that Player 1 assigns to Heads with $\sigma_{1T} = (1 - \sigma_{1H})$ be the probability that Player 1 assigns to Tails. In order to make Player 2 indifferent among his pure strategies, we need $E_2[Heads] = E_2[Tails]$. The expected values for Player 2 of playing Heads and Tails are:

$$\begin{aligned} E_2[Heads] &= \sigma_{1H} * (-1) + (1 - \sigma_{1H}) * 1 \\ E_2[Tails] &= \sigma_{1H} * 1 + (1 - \sigma_{1H}) * (-1) \end{aligned}$$

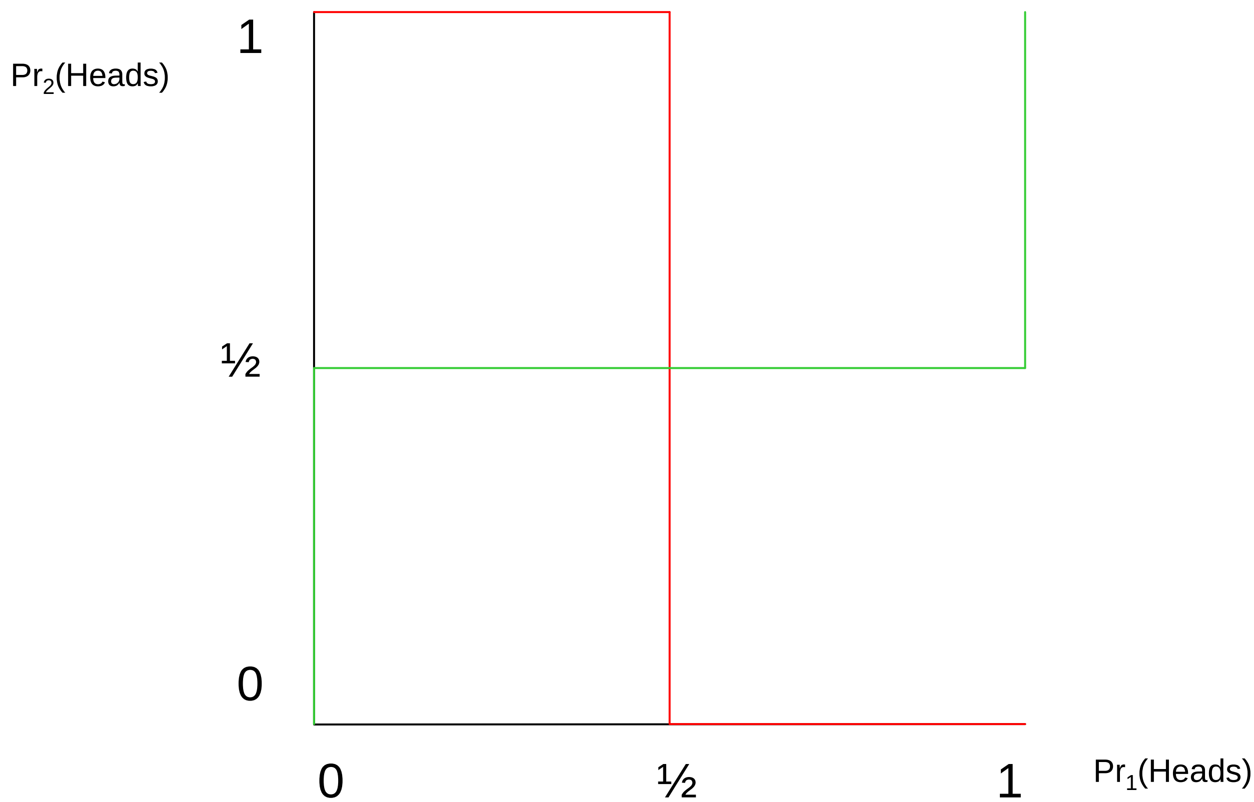

Figure 2: Best response functions for the Matching Pennies game.

Now, set these 2 equal and solve for σ_{1H}.

$$\begin{aligned}\sigma_{1H}*(-1)+(1-\sigma_{1H})*1 &= \sigma_{1H}*1+(1-\sigma_{1H})*(-1)\\ -\sigma_{1H}+1-\sigma_{1H} &= \sigma_{1H}-1+\sigma_{1H}\\ 1-2\sigma_{1H} &= 2\sigma_{1H}-1\\ 2 &= 4\sigma_{1H}\\ \frac{2}{4} &= \sigma_{1H}\end{aligned}$$

We can then show that $\sigma_{1T}=\frac{2}{4}$ as well. A similar process will provide $\sigma_{2H}=\sigma_{2T}=\frac{1}{2}$. So the Nash Equilibrium to the Matching Pennies game is Player 1 chooses Heads with probability $\frac{1}{2}$ and Tails with probability $\frac{1}{2}$ and Player 2 chooses Heads with probability $\frac{1}{2}$ and Tails with probability $\frac{1}{2}$. It does not have to be the case that the mixed strategy Nash Equilibrium is symmetric, nor is it the case that the existence of pure strategy Nash Equilibrium will eliminate the possibility of a mixed strategy Nash Equilibrium. Consider the coordination game. Player 1 and Player 2 have 2 locations at which they can meet, Picasso's and BBs. However, they are unable to communicate on where to meet (you have to realize that this game was created prior to the popularity of cell phones). They prefer meeting to not meeting, but Player 1 prefers meeting at Picasso's to meeting at BBs and Player 2 prefers meeting at BBs to meeting at Picasso's. The matrix representation is below:

		Player 2	
		Picasso's	BBs
Player 1	Picasso's	$3,2$	$0,0$
	BBs	$0,0$	$2,3$

In this game it is easy to see that there are two pure strategy Nash Equilibria. One is Player 1 chooses Picasso's and Player 2 chooses Picasso's and the other is Player 1 chooses BBs and Player 2 chooses BBs. However, there is also a mixed strategy Nash Equilibrium to this game, where Player 1 chooses Picasso's 60% of the time and BBs 40% of the time and Player 2 chooses Picasso's 40% of the time and BBs 60% of the time. You should check to make sure that this is indeed a Nash Equilibrium of the game.

2.2.1 Existence

Recall that when we studied consumer and producer theory we were concerned with existence and uniqueness of the consumer and producer problems, as well as existence and uniqueness of competitive equilibrium. It is no different with Nash Equilibrium. How do we know that a solution to the game actually exists?

Proposition 12 *Every game $\Gamma_N=[I,\{\Delta(S_i)\},\{u_i(\cdot)\}]$ in which the sets $S_1,...,S_I$ have a finite number of elements has a mixed strategy Nash Equilibrium (note that it could be degenerate, like the Prisoner's Dilemma).*

So, if we allow for mixed strategies and the players have a finite number of strategies to choose from then we are guaranteed to have at least one Nash Equilibrium.

Proposition 13 *A Nash Equilibrium exists in game $\Gamma_N=[I,\{S_i\},\{u_i(\cdot)\}]$ if for all $i=1,...,I$*

1. *S_i is a nonempty, convex, and compact subset of some Euclidean space $\mathbb{R}^M$*
2. *$u_i(s_1,...,s_I)$ is continuous in $(s_1,...,s_I)$ and quasiconcave in s_i*

If we make some restrictions about the strategy set S_i then we can show that an equilibrium in pure strategies exists. These results hinge on fixed-point theorems. The particular fixed-point theorem used depends on the assumptions one makes about the strategy sets and payoff function. The assumptions in the proposition are for Brouwer's fixed-point theorem. A fixed-point theorem basically says that there is a point in the set that maps back to itself. In the case of the games we are playing, there are points in the best response correspondences of players that map back into themselves.

3 Sequential Move Games

Thus far we have examined games in which players make moves simultaneously (or without observing what the other player has done). Using the normal (strategic) form representation of a game we can identify sets of strategies that are best responses to each other (Nash Equilibria). We now focus on sequential games of complete information. We can still use the normal form representation to identify NE but sequential games are richer than that because some players observe other players' decisions before they take action. The fact that some actions are observable may cause some NE of the normal form representation to be inconsistent with what one might think a player would do.

Here's a simple game between an Entrant and an Incumbent. The Entrant moves first and the Incumbent observes the Entrant's action and then gets to make a choice. The Entrant has to decide whether or not he will enter a market or not. Thus, the Entrant's two strategies are "Enter" or "Stay Out". If the Entrant chooses "Stay Out" then the game ends. The payoffs for the Entrant and Incumbent will be 0 and 2 respectively. If the Entrant chooses "Enter" then the Incumbent gets to choose whether or not he will "Fight" or "Accommodate" entry. If the Incumbent chooses "Fight" then the Entrant receives -3 and the Incumbent receives -1. If the Incumbent chooses "Accommodate" then the Entrant receives 2 and the Incumbent receives 1. This game in normal form is

		Incumbent	
		Fight if Enter	Accommodate if Enter
Entrant	Enter	$-3, -1$	$2, 1$
	Stay Out	$0, 2$	$0, 2$

Note that there are two pure strategy Nash Equilibria (PSNE) to this game. One is that the Entrant chooses Enter and the Incumbent chooses Accommodate and the other is that the Entrant chooses Stay Out and the Incumbent chooses Fight.[2] Of these 2 PSNE, which seems more "believable"? The NE where the Entrant chooses Stay Out and the Incumbent chooses Fight is only a NE if the Entrant thinks that the Incumbent's choice of Fight is credible. But what does the Entrant know? The Entrant knows that if he chooses Enter that the Incumbent will not choose Fight but will choose Accommodate. If you are the Entrant, are you worried about the Incumbent choosing Fight (as this game is structured)? No, because if you choose Enter, the best thing for the Incumbent to do for himself at that point would be to choose Accommodate. Thus, we can rule out the NE of Stay Out, Fight because the choice of Fight is not credible (Note: this does NOT mean that Stay Out, Fight is NOT a NE, it just means that it relies on a noncredible threat).

3.1 Extensive Form Representation

We can represent the sequential game using the extensive form representation or game tree. The extensive form representation, Γ_E, has more components than the normal form representation. Recall that the normal form representation required that we only need to know how many players there were, which strategies were available to each player, and which payoffs occurred as a result of the players' strategy choices. With an extensive form game we also need to consider the fact that players move at different points in time. An extensive form game will consist of the following basic items: players, decision nodes, information sets, strategies, and payoffs.[3] Note that the components of a normal form game are all here, so that any extensive form game may also be represented as a normal form game. The additional items are decision nodes and information sets. A decision node is a point where a player makes a decision. An information set is what the player knows at a certain node. The game begins with a single decision node, called the initial node. The player who makes a decision at a node is listed and the strategies available to that player are then drawn as lines from that node (the branches of the game tree). If the game ends after a player makes a decision, then the players reach a terminal node and payoffs are listed. The payoffs correspond to the path played out by the strategy choices that lead to the terminal node. If the game continues after a player makes a decision then the players reach another decision node. A different player will be able to take an action, and his strategies will be represented as branches extending from that node. This process continues until the

[2] To be complete, there is also a mixed strategy Nash Equilibrium (MSNE) where the Entrant chooses Stay Out with probability 1 and the Incumbent chooses Fight with probability $\frac{2}{5}$ and Accommodate with probability $\frac{3}{5}$.

[3] There is a much more formal definition on page 227 of the text.

game reaches a terminal node after strategy choices are made. The game tree representation of the Entrant, Incumbent game is in Figure 3.1. Note that an actual game tree does not include all the labels, but I have included them for reference.

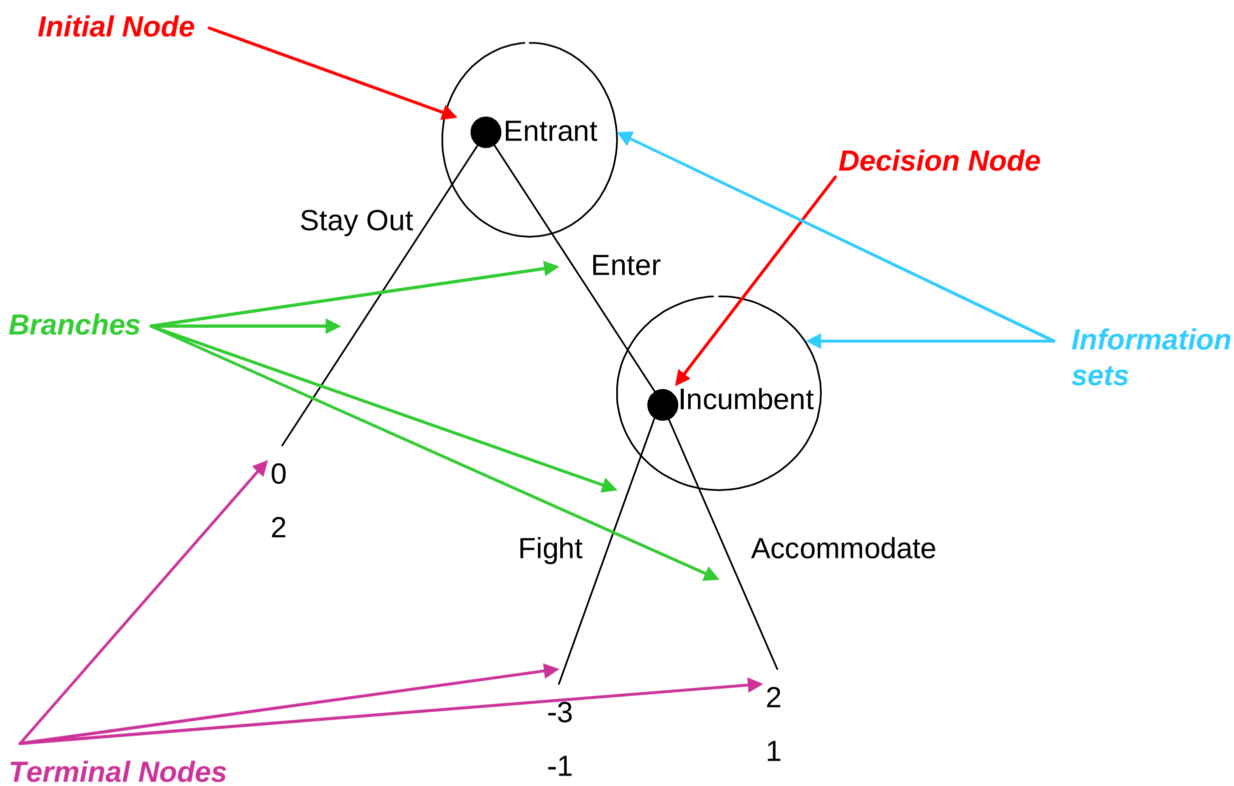

Game tree with its components labeled.

The actual game tree, without the labels, would look like Figure 3.1.

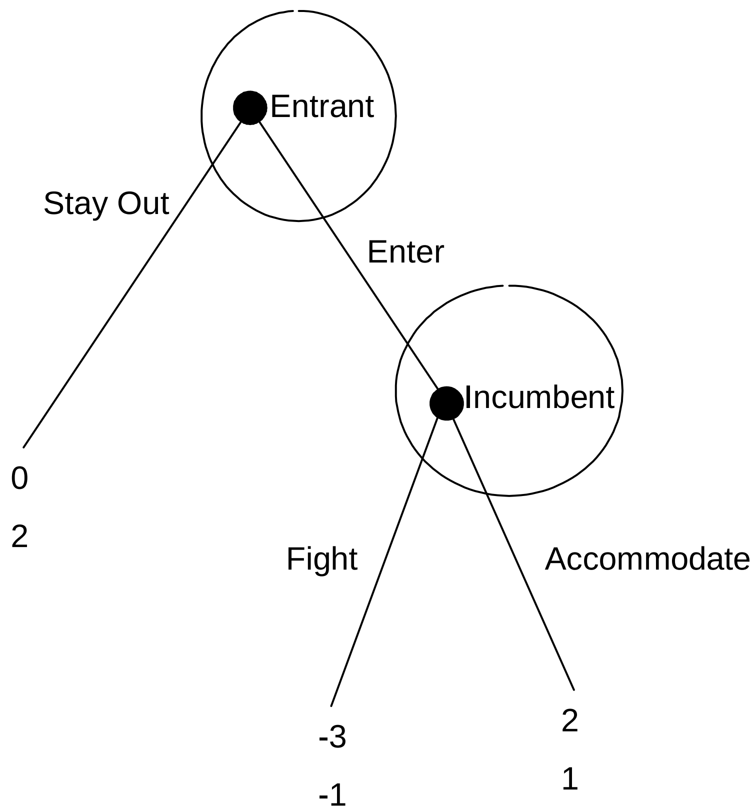

Game tree without the components labeled.

Much simpler, but note that the players, strategies, and payoffs are still listed. You might ask why we circle the node and label it an information set. It is possible that players do not know which node they are at, so that the player's information set contains multiple nodes. Thus, the information set would be a circle around both nodes. Consider the simultaneous move Prisoner's Dilemma game. Prisoner 2 does not know which choice is made by Prisoner 1, so his information set contains both nodes, as in Figure 3.Contrast this with the sequential version of the Prisoner's Dilemma game where Prisoner 2 knows what Prisoner 1 has chosen in Figure 4.The games look similar, but they are slightly different as we will discuss shortly.

There is one other small detail in extensive form games. It may be that one or more players has an infinite number of strategies. Consider a game in which players may choose to produce any quantity of an item greater than or equal to 0 or less than or equal to the quantity consumers would purchase when price equals 0. The strategy space is then any real number in the interval $[0, Q_0]$, where Q_0 is the quantity consumers would purchase when price equals 0. Since we cannot represent the strategies with a finite number of branches, we would use a dashed line between two branches to represent an interval. The branches would be labeled 0 and Q_0. If the other player does not know the choice of the first player (the game is simultaneous) then both nodes extending from the two branches as well as all the nodes represented by the dashed line are in the information set, so we circle both nodes and the dashed line. If the second player does observe the first player's choice, then we simply draw a circle around some elements of the dashed line but not the two nodes drawn from the branches.

Representing an extensive form game in normal form We know that with a normal form game we only need to know the players, strategies, and payoffs. Representing Entry, Incumbent in normal form is easy: since there are only 2 players and 2 strategies we only need a 2x2 matrix. We have already seen the representation of the simultaneous Prisoner's Dilemma. But what does the normal form representation of the sequential Prisoner's Dilemma look like? There are 2 players, Prisoner 1 and Prisoner 2. Prisoner 1 has 2 strategies Confess and Not Confess. So far everything looks the same. But Prisoner 2 now has 4 strategies? How can that be? He only chooses Confess or Not Confess, doesn't he? While that is

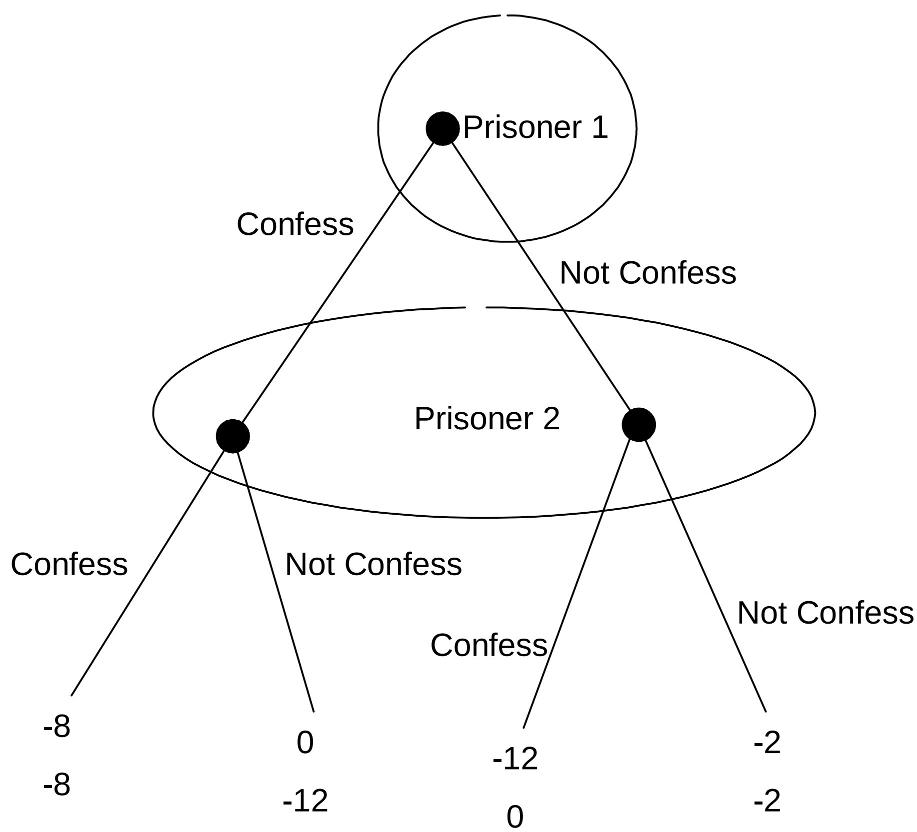

Figure 3: Simultaneous move Prisoner's Dilemma game.

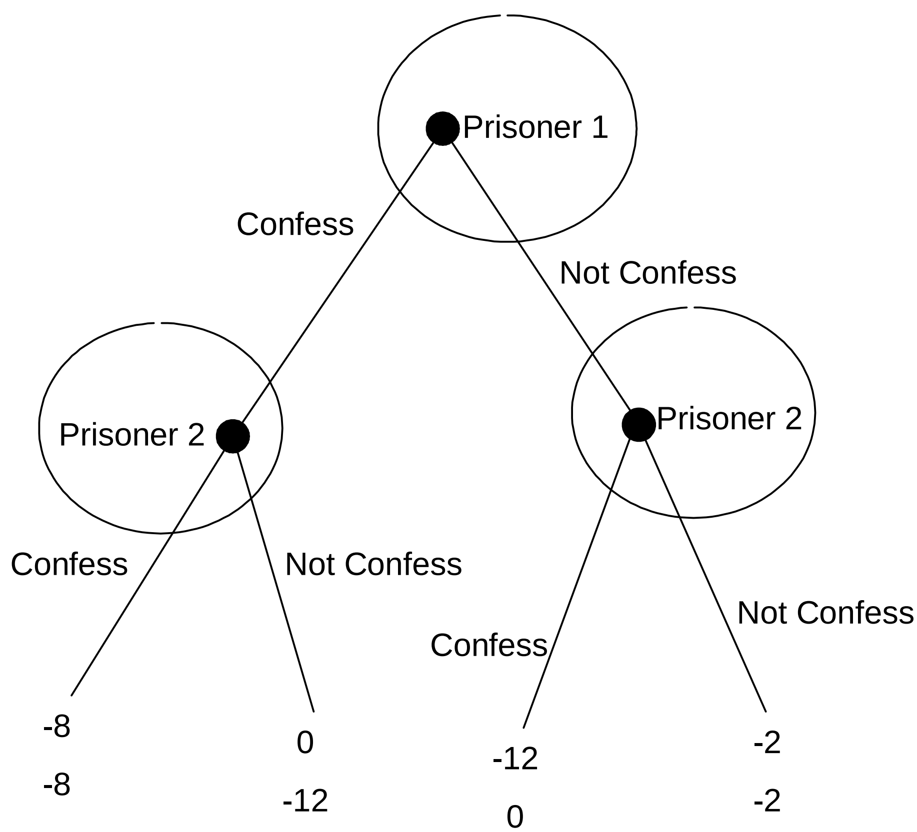

Figure 4: Sequential move Prisoner's Dilemma game.

true, Prisoner 2 now has 2 information sets – he needs to specify an action at EVERY information set. So Prisoner 2's 4 strategies are:

1. Confess if Prisoner 1 confesses, Confess if Prisoner 1 does not confess (essentially Always Confess)
2. Confess if Prisoner 1 confesses, Not Confess if Prisoner 1 does not confess (essentially play the same as Prisoner 1)
3. Not Confess if Prisoner 1 confesses, Confess if Prisoner 1 does not confess (essentially play the opposite of Prisoner 1)
4. Not Confess if Prisoner 1 confesses, Not Confess if Prisoner 1 does not confess (essentially Always Not Confess)

Given that Prisoner 2 has 4 strategies, we now have a 2x4 (or a 4x2) matrix representation of the sequential Prisoner's Dilemma.

		Prisoner 1	
		Confess (C)	Not Confess (NC)
	Confess if P1 C, Confess if P1 NC	$-8,-8$	$0,-12$
Prisoner 2	Confess if P1 C, Not Confess if P1 NC	$-8,-8$	$-2,-2$
	Not Confess if P1 C, Confess if P1 NC	$-12,0$	$0,-12$
	Not Confess if P1 C, Not Confess if P1 NC	$-12,0$	$-2,-2$

This normal form representation illustrates what a strategy is for a player in a sequential game. Finding the NE, we see that the only pure strategy NE to the sequential game is that Prisoner 1 chooses Confess and Prisoner 2 chooses "Confess if P1 C, Confess if P1 NC" or "Always Confess". Thus, there are no NE that rely on noncredible threats in this game, so there is no difference in the OUTCOME in the sequential and the simultaneous Prisoner's Dilemma games. But the NE are NOT the same because Prisoner 2 has a different strategy set in the two games.

3.2 Finding "credible" NE in an extensive form game

Since we can represent any extensive form game in strategic form, why bother with the game tree? The game tree allows us to use a concept called backward induction to find those NE which are sequentially rational and eliminate those NE which are not. In the Entrant, Incumbent game the NE of Stay Out, Fight is NOT sequentially rational while the NE of Enter, Accommodate is. Sequentially rational means that all players are choosing optimally at any point in the tree. Backward induction says that to find theses sequentially rational NE one starts at the end of the tree (the terminal nodes) and works backwards, choosing optimally at each decision node and "eliminating" the branches of the tree that are not chosen. The NE found by this method is known as the subgame perfect Nash Equilibrium (SPNE – do not confuse with pure strategy Nash Equilibrium, PSNE). What is a subgame?

Definition 14 *A subgame of an extensive form game Γ_E is a subset of the game having the following properties*

1. *it begins with an information set containing a single decision node, contains all the decision nodes that are successor nodes of this node, and contains only those nodes*
2. *If decision node x is in the subgame, then every $x' \in H(x)$ is also, where $H(x)$ is the information set that contains x (there are no broken information sets)*

Now that we know what a subgame is, we can define a SPNE. Note that the entire game tree is a subgame.

Definition 15 *A profile of strategies $\sigma = (\sigma_1, ..., \sigma_I)$ in an I-player extensive form game Γ_E is a subgame perfect Nash Equilibrium if it induces a Nash Equilibrium in every subgame of Γ_E*

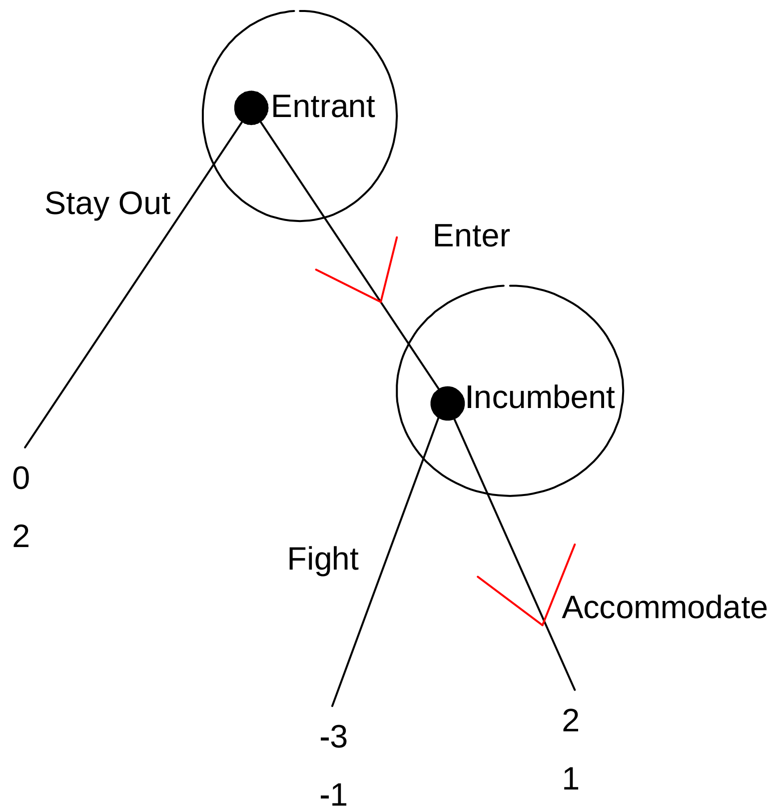

Figure 5: Illustration of SPNE of the Entrant, Incumbent game.

Basically, if we look at each of the subgames in a game tree we want players to be playing a NE in each of the subgames. This is why backward induction works as a solution technique. We start at the smallest subgames and work towards the largest subgame (the entire game) finding optimal choices at each of the smaller subgames. Consider the Entrant, Incumbent game again. We already know that there are 3 NE (2 PSNE and 1 MSNE). However, there is only 1 subgame perfect NE (SPNE), which is the Entrant Enters and the Incumbent Accommodates. To see this, start at the smallest subgame, which is the Incumbent's decision node and the two branches extending from it. A rational Incumbent would choose to Accommodate in this subgame because $1 > -1$. Knowing this, the Entrant can now eliminate the "Fight" branch from the tree, because a rational Incumbent will not play this strategy if the game gets to this point. The Entrant now has to choose between Stay Out, which has a payoff of 0 for the Entrant, and Enter, which results in a payoff of 2 for the Entrant. Since $2 > 0$, the Entrant will choose Enter.The difference between Figure 5 and Figure 3.1 is that the subgame perfect choices made by the players now have red arrows indicating the choices the players would make. This is standard notation for indicating choices in an extensive form game (the arrows at least, not necessarily the red color). We can also find the SPNE of the sequential prisoner's dilemma game using the same techniques, and it will show that Prisoner 1 Confesses, Prisoner 2 Always Confesses is the SPNE of that game. Note that all SPNE are NE, but not that all NE are SPNE. As for existence of NE and SPNE in sequential games, we have two propositions.

Proposition 16 *(Zermelo's Theorem) Every finite game of perfect information* Γ_E *has a pure strategy Nash Equilibrium that can be derived through backward induction. Moreover, if no player has the same payoffs at any two terminal nodes, then there is a unique Nash Equilibrium that can be derived in this manner.*

Proposition 17 *Every finite game of perfect information* Γ_E *has a pure strategy subgame perfect Nash Equilibrium that can be derived through backward induction. Moreover, if no player has the same payoffs at* *any two terminal nodes, then there is a unique subgame perfect Nash Equilibrium that can be derived in this manner.*

Note that we need perfect information in order to have these propositions hold, not just common knowledge. Perfect information requires that all information sets contain a single decision node. Thus, the simultaneous move Matching Pennies game and the simultaneous move Prisoner's Dilemma game do not have perfect information, so there is no guarantee that there is a PSNE to either of those games (there might be though).

4 Repeated Interactions

We have studied simultaneous and sequential games, and those games have essentially been one-shot games in nature. One-shot games are a starting point for the discussion of games, but many "games" are played repeatedly between players. Consider the following game:

		Player 2	
		Defect	Cooperate
Player 1	Defect	8, 8	32, 4
	Cooperate	4, 32	25, 25

Note that Defect is a dominant strategy for both players, but that if they could agree to Cooperate they would earn more. Note that this game is simply a Prisoner's Dilemma without the negative payoffs and the underlying story. What if the players played this game multiple, but finite, times? What would the SPNE of the repeated game be? Consider two repetitions of the game, and recall that a SPNE must induce NE in every subgame. Start with the ending subgames. Since these are simply Prisoner's Dilemmas the players must both choose Defect in order to induce NE of the subgame. Now the players have a choice of receiving either 16 and 40 (if they choose Defect) or 12 and 25 (if they choose Cooperate). Since $16 > 12$ and $40 > 25$ the players will both choose Defect in the initial play of the game. Thus, the SPNE of this twice repeated Prisoner's Dilemma is for player 1 to always choose Defect at any decision node and for Player 2 to always choose Defect at any decision node. Regardless of how many times this game is *finitely* repeated, the only SPNE of this game will be for both of the players to choose Defect at any decision node. Thus, any attempt at cooperation in a finitely repeated Prisoner's Dilemma should unravel according to SPNE.

One might think that the players should be able to cooperate if they are playing this game repeatedly. After all, if the game is played 1000 times it is a lot better to receive 25 each period than it is to receive 8 each period. But SPNE is what it is for this game. However, what if the game was repeated *infinitely*? The first question to ask is if infinite repetition even makes sense given that the lives of humans are finite (at least to the best of our knowledge). Consider an "economic agent" that is a corporation. The corporation may be infinitely lived as it passes from one owner to the next. Second, while human lives may be finite there is (usually) some uncertainty as to when one's life will end. We can show that having an uncertain endpoint is consistent with infinite repetition of a game. Finally, and this answers the question of why infinite repetition before discussing the concept, people DO cooperate with one another on a daily basis. Infinite repetition of a one-shot game like the Prisoner's Dilemma will allow the (Cooperate, Cooperate) outcome to occur at every repetition of the game as part of a viable SPNE of the game. Be warned, however, that infinite repetition of the one-shot Prisoner's Dilemma also allows the (Defect, Defect) outcome to occur at every repetition of the game as part of a viable SPNE of the game. Thus, while infinite repetition will allow cooperation as part of the SPNE, infinite repetition also allows for a multiplicity of equilibria. This is what is known as the "embarrassment of riches" of infinitely repeated games. Recall from earlier discussions that economists like to answer two questions when discussing the concept of equilibrium: Does an equilibrium exist and is it unique? We focus on showing the sufficient conditions for equilibrium to exist and skip the uniqueness question when discussing infinitely repeated games. There is a third question which some of you may be interested in that we will not discuss in detail: Among the multiple equilibria that exist, how do the players choose one of them? This is essentially the basic question of evolutionary economics, which attempts to move economics from a physics framework to a biological one. Well this seems novel, it has roots dating back at least to Alfred Marshall, who wrote the primary economics text (Principles of Economics – not a very clever title, but to the point) around the turn of the century – the 20^{th} century (published in 1890).

4.1 Evaluating strategies in infinite games

In order to evaluate strategies in infinite games it will be necessary to add a particular parameter to the discussion. The parameter added will be the player's discount rate, δ. It is assumed that $\delta \in [0, 1)$, and that players have exponential discounting. All that exponential discounting means is that a payoff one time period from today is discounted at δ and a payoff two time periods from today is discounted at δ^2, etc. Thus, a player's payoff stream from the infinite game would look like:

$$\delta^0 \Pi_0 + \delta^1 \Pi_1 + \delta^2 \Pi_2 + \delta^3 \Pi_3 + ...$$

where Π_k denotes the player's payoff in each period k. The $\delta \in [0, 1)$ assumption will be justified shortly.[4] It is typically assumed that players (and people in general) prefer \$1 today to \$1 tomorrow, and \$1 tomorrow to \$1 two days from now. Thus, the sooner a player receives a payoff the less he discounts it. Why add this discount rate? Well, if we do not have a discount rate then the players' payoffs from following ANY strategy (assuming that there are no negative payoffs that the player could incur) of an infinite game would be infinite. Well, that's not very interesting. This is also why we assume that $\delta < 1$ rather than $\delta \leq 1$. If $\delta = 1$, then a player weights all payoffs equally regardless of the time period, and this leads to an infinite payoff. If $\delta = 0$, then the player will only care about the current period. As δ moves closer to 1, the player places more weight on future periods. It is possible to motivate this discount rate from a present value context, which I believe would make $\delta = \frac{1}{1+r}$, where r is "the interest rate". Thus, if $r = 0.05$, then $\delta \approx 0.95$. All this says is that getting \$1 one period from today is like getting 95 cents today, and getting \$1 two periods from today is like getting 90.7 cents today. While this interpretation of the discount rate is the most closely linked to economic behavior, we will not assume that the discount rate is directly related to the interest rate, but that it is simply a parameter that states how players value payoffs over time.

Now, suppose that players 1 and 2 use the following strategies:

Player 1 chooses Cooperate in the initial period (at time $t = 0$) and continues to choose Cooperate at every decision node unless he observes that player 2 has chosen Defect. If Player 1 ever observes Player 2 choosing Defect then Player 1 will choose Defect at every decision node after that defection. Player 2's strategy is the same. These strategies call for Cooperation at every decision node until a Defection is observed and then Defection at every decision node after Defection is observed. Note that this is a *potential* SPNE because it is a set of strategies that specifies an action at every decision node of the game. The question then becomes whether or not this is a SPNE of the game. Recall that a strategy profile is an SPNE if and only if it specifies a NE at every subgame. Although each subgame of this game has a distinct history of play, all subgames have an identical structure. Each subgame is an infinite Prisoner's Dilemma exactly like the game as a whole. To show that these strategies are SPNE, we must show that after any previous history of play the strategies specified for the remainder of the game are NE.

Consider the following two possibilities:

1. A subgame that contains a deviation from the Cooperate, Cooperate outcome somewhere prior to the play of the subgame

2. A subgame that does not contain a deviation from the Cooperate, Cooperate outcome

If a subgame contains a deviation then the players will both choose Defect, Defect for the remainder of the game. Since this is the NE to the one-shot version (or stage game) of the Prisoner's Dilemma, it induces a NE at every subgame. Thus, the "Defect if defection has been observed" portion of the suggested strategy induces NE at every subgame.

Now, for the more difficult part. Suppose that the players are at a subgame where no previous defection has occurred. Consider the potential of deviation from the proposed strategy in period $\tau \geq t$, where t is the current period. If player 2 chooses Defect in period τ he will earn $\delta^\tau \Pi^{Deviate} + \delta^\tau \sum_{i=1}^{\infty} \delta^i \Pi^D$ for the remainder of the game, where $\Pi^{Deviate}$ is player 2's payoff from deviating and Π^D is his payoff each period from the (Defect, Defect) outcome. If player 2 chooses to follow the proposed strategy, then he will earn $\delta^\tau \sum_{i=0}^{\infty} \delta^i \Pi^C$, where Π^C is his payoff from the (Cooperate, Cooperate) outcome. The question

[4] The exponential discounting assumption is used because it allows for time consistent preferences. Hyperbolic discounting is another type of discounting that has been suggested as consistent with choices made by individuals in experiments, although hyperbolic discounting does not necessarily lead to time consistent preferences.

then becomes under what conditions will the payoff from deviating be greater than that from the payoff of following the proposed strategy. To find the condition simply set up the inequality:

$$\delta^\tau \Pi^{Deviate} + \delta^\tau \sum_{i=1}^{\infty} \delta^i \Pi^D \geq \delta^\tau \sum_{i=0}^{\infty} \delta^i \Pi^C$$

We can cancel out the δ^τ terms to obtain:[5]

$$\Pi^{Deviate} + \sum_{i=1}^{\infty} \delta^i \Pi^D \geq \sum_{i=0}^{\infty} \delta^i \Pi^C$$

Now, using results on series from Calculus, we have:

$$\Pi^{Deviate} + \frac{\delta}{1-\delta}\Pi^D \geq \frac{1}{1-\delta}\Pi^C$$

Now, we can substitute in for $\Pi^{Deviate}, \Pi^D$, and Π^C from our game to find:

$$32 + 8\frac{\delta}{1-\delta} \geq 25\frac{1}{1-\delta}$$

Or:

$$\begin{aligned} 32 - 32\delta + 8\delta &\geq 25 \\ 7 - 24\delta &\geq 0 \\ 7 &\geq 24\delta \\ \frac{7}{24} &\geq \delta \end{aligned}$$

Thus, choosing to deviate from the proposed strategy only provides a higher payoff if $\delta \leq \frac{7}{24}$, so that continuing to cooperate is a best response if $\delta \geq \frac{7}{24}$. The discount rate will be a key factor in determining whether or not a proposed equilibrium is a SPNE. In fact, when looking at infinitely repeated games, it is best to have a particular strategy in mind and then check to see what the necessary conditions are for it to be a SPNE, given the multiplicity of equilibria.

Are there other SPNE to the game? Well, consider a modified version of the game:

		Player 2	
		Defect	Cooperate
Player 1	Defect	8, 8	80, 4
	Cooperate	4, 80	25, 25

The only change in this game is that the payoff of 32 that the player received from Defecting when the other player Cooperates has been changed to 80. We can show that both players using a strategy of cooperating until a defection occurs (the same proposed strategy from before) is a SPNE if:

$$80 + 8\frac{\delta}{1-\delta} \geq 25\frac{1}{1-\delta}$$

or $\delta \geq \frac{55}{72}$. Thus, if both players are sufficiently patient then the proposed strategy is still a SPNE. Note that the discount rate increased in this example since the payoff to deviating increased. But, is there a strategy that yields higher payoffs? What if the following strategies were used by players 1 and 2:

If no deviation has occurred, Player 1 chooses Defect in all even time periods and chooses Cooperate in all odd time periods. If a deviation occurs Player 1 always chooses Defect.

If no deviation has occurred, Player 2 chooses Cooperate in all even time periods and chooses Defect in all odd time periods. If a deviation occurs Player 2 always chooses Defect.

A deviation (from player 1's perspective) occurs when Player 2 chooses Defect in an even time period. A deviation (from player 2's perspective) occurs when Player 1 chooses Defect in an odd time period. Note that we start the game at time $t = 0$, so that Player 1 receives 80 first.

[5]This canceling out of the δ^τ terms typically leads to the assumption that if deviation is going to occur in an infinitely repeated game it will occur in the first time period. I proceed under this assumption in later examples.

Look at what this strategy would do. It would cause the outcome of the game to alternate between the $(Defect, Cooperate)$ and $(Cooperate, Defect)$ outcomes, giving the players alternating periods of payoffs of 80 and 4, as opposed to 25 each period using the "cooperate until defect is observed, then always defect" strategy. On average (and ignoring discounting for a moment), each player would receive 42 per period under this new strategy and only 25 per period under the old. Is the new strategy a SPNE? We should check for both players now that they are receiving different amounts of payoffs in different periods.

For Player 1:

$$\begin{aligned} \Pi^{Deviate} &= 80 + \textstyle\sum_{i=1}^{\infty} \delta^i 8 \\ \Pi^{C} &= \textstyle\sum_{i=0}^{\infty} \delta^{2i} 80 + \sum_{i=0}^{\infty} \delta^{2i+1} 4 \end{aligned}$$

If $\Pi^C \geq \Pi^{Deviate}$ then Player 1 will choose NOT to deviate:

$$\begin{aligned} 80\frac{1}{1-\delta^2} + 4\frac{\delta}{1-\delta^2} &\geq 80 + 8\frac{\delta}{1-\delta} \\ 80 + 4\delta &\geq 80\left(1-\delta^2\right) + 8\delta\left(1+\delta\right) \\ 4\delta &\geq -80\delta^2 + 8\delta + 8\delta^2 \\ 72\delta^2 - 4\delta &\geq 0 \\ 18\delta - 1 &\geq 0 \\ \delta &\geq \frac{1}{18} \end{aligned}$$

This is true, for any $\delta \geq \frac{1}{18}$.

For Player 2:

$$\begin{aligned} \Pi^{Deviate} &= \textstyle\sum_{i=0}^{\infty} \delta^i 8 \\ \Pi^{C} &= \textstyle\sum_{i=0}^{\infty} \delta^{2i} 4 + \sum_{i=0}^{\infty} \delta^{2i+1} 80 \end{aligned}$$

If $\Pi^C \geq \Pi^{Deviate}$ then Player 2 will choose NOT to deviate:

$$\begin{aligned} 4\frac{1}{1-\delta^2} + 80\frac{\delta}{1-\delta^2} &\geq 8\frac{1}{1-\delta} \\ 4 + 80\delta &\geq 8 + 8\delta \\ 72\delta &\geq 4 \\ \delta &\geq \frac{1}{18} \end{aligned}$$

Thus, both players need to have a discount rate greater than or equal to $\frac{1}{18}$ to support this strategy. Note that this discount rate is much lower than the one needed to support the "cooperate until defect is observed, then always defect" strategy. However, it also illustrates the "embarrassment of riches" of infinitely repeated games.

4.2 Some formalities

We will now formalize some of these concepts. The focus is on 2-player games. In the one-period stage game, each player i has a compact strategy set S_i, where $q_i \in S_i$ is a particular feasible action for player i.

Let $q = (q_1, q_2)$ and $S = S_1 x S_2$.

Let $\pi_i(q_i, q_j)$ be player i's payoff function.

Let $\widehat{\pi}_i(q_j) = Max_{q \in S_i} \pi_i(q_i, q_j)$ be player i's one period best response payoff given that his rival chooses q_j.

Let $q^* = (q_1^*, q_1^*)$ denote the unique PSNE to the one-period stage game (a simplifying assumption).

A pure strategy in this game for player i, s_i, is a sequence of functions, $\{s_{it}(\cdot)\}_{t=1}^{\infty}$ mapping from the history of previous action choices (denoted H_{t-1}) to his action choice in period t, $s_{it}(H_{t-1}) \in S_i$.

The set of all pure strategies for player i is denoted Σ_i, and $s = (s_1, s_2) \in \Sigma_1 x \Sigma_2$ is a profile of pure strategies for the players.

Any pure strategy profile $s = (s_1, s_2)$ induces an outcome path $Q(s)$, which is an infinite sequence of actions $\{q_t = (q_{1t}, q_{2t})\}_{t=1}^{\infty}$ that will actually be played when the players follow strategies s_1 and s_2.

Player i's discounted payoff from outcome path Q is denoted by $v_i(Q) = \sum_{t=0}^{\infty} \delta^t \pi_i(q_{1+t})$.

Player i's average payoff from outcome path Q is $(1-\delta) v_i(Q)$.

Player i's discounted continuation payoff from some point t onward is $v_i(Q,t) = \sum_{\tau=0}^{\infty} \delta^{\tau} \pi_i(q_{t+\tau})$.

We already know that the strategies that call for player i to play his stage game NE q_i^* in every period, regardless of the prior history, constitute an SPNE for any $\delta < 1$.

4.2.1 Nash reversion and the Nash reversion Folk Theorem

Nash reversion is essentially the "punishment" we have been discussing – if one player fails to "cooperate", the other player "punishes" by reverting to the stage game NE. It was well-known that this was a solution to the infinitely repeated game before someone decided to write it down, hence the term "Folk Theorem".

Definition 18 *A strategy profile $s = (s_1, s_2)$ in an infinitely repeated game is one of Nash reversion if each player's strategy calls for playing some outcome path Q until someone deviates and playing the stage game NE $q^* = (q_1^*, q_2^*)$ thereafter.*

Lemma 19 *A Nash reversion strategy profile that calls for playing outcome path $Q = \{q_{1t}, q_{2t}\}_{t=1}^{\infty}$ prior to any deviation is a SPNE if and only if*

$$\widehat{\pi}_i(q_{jt}) + \frac{\delta}{1-\delta}\pi_i\left(q_i^*, q_j^*\right) \leq v_i(Q,t)$$

where $j \neq i$ for all t and $i = 1,2$.

This formalizes what we have already been discussing in the context of the Prisoner's Dilemma game.

Proposition 20 *Consider an infinitely repeated game with $\delta > 0$ and $S_i \subset \mathbb{R}$ for $i = 1,2$. Suppose also that $\pi_i(q)$ is differentiable at $q^* = (q_1^*, q_2^*)$ with $\partial \pi_i(q_1^*, q_2^*)/\partial q_j \neq 0$ for $j \neq i$ and $i = 1,2$. Then there is some $q' = (q_1', q_2')$ with $[\pi_1(q_1', q_2'), \pi_2(q_1', q_2')] >> [\pi_1(q_1^*, q_2^*), \pi_2(q_1^*, q_2^*)]$ whose infinite repetition is the outcome path of an SPNE that uses Nash reversion.*

This proposition states that with continuous strategy sets and differentiable payoff functions, as long as there is some possibility for joint improvement in payoffs around the stage game NE some cooperation can be sustained.

Proposition 21 *Suppose that outcome path Q can be sustained as an SPNE outcome path using Nash reversion when the discount rate is δ. Then it can be so sustained for any $\delta' \geq \delta$.*

Hopefully that proposition is self-explanatory ...

Proposition 22 *For any pair of actions $q = (q_1, q_2)$ such that $\pi_i(q_1, q_2) > \pi_i(q_1^*, q_2^*)$ for $i = 1,2$ there exists a $\underline{\delta} < 1$ such that, for all $\delta > \underline{\delta}$, infinite repetition of $q = (q_1, q_2)$ is the outcome path of an SPNE using Nash reversion strategies.*

This is the important proposition as it states that any stationary outcome path that gives each player a discounted payoff that exceeds the payoff arising from infinite repetition of the stage game NE can be sustained as an SPNE if δ is sufficiently close to 1. Note that this proposition only applies to stationary paths. It is possible to extend the argument to include non-stationary paths with average payoff vectors greater than the stage game NE.

The text also discusses some other propositions which show that payoff vectors LESS than the stage game NE payoff vector can be supported using punishment strategies that are harsher than Nash reversion. Essentially, Player 1 threatens to punish Player 2 by forcing Player 2 to accept the minimum amount he possibly could be forced to accept. In the simple Prisoner's Dilemma we have been discussing this is the same as Nash reversion because Player 2 could not be forced to accept anything less than the stage game NE.