Simultaneous Move Games

These notes essentially correspond to parts of chapters 7 and 8 of Mas-Colell, Whinston, and Green.

1 Introduction

Up to now in the Micro Theory sequence you have typically 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. 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; 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* $\mathcal{H}_i$ *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* $s_i : \mathcal{H}_i \rightarrow \tilde{A}$ *such that* $s_i(H) \in C(H)$ *for all* $H \in \mathcal{H}_i$*.*

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.

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 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 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

		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.

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".

2.1 Finding Nash equilibria – IEDS

Now that we have a definition of Nash equilibrium, how do we determine which sets of strategies constitute Nash equilibria? One method, which could be very time consuming (particularly if done by hand – possibly less so if done by computer), would be to determine every possible combination of strategies in a game and then determine if they satisfy the definition we have of Nash equilibrium. Fortunately, there are some shortcuts. Here is an example:

		Player 2		
		Left	Center	Right
	Top	7, 4	6, 3	4, 11
Player 1	Middle	8, 8	10, 4	6, 7
	Bottom	18, 7	11, 9	4, 6

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

In the Prisoner's Dilemma both players had a strictly dominant strategy of Confess, which made finding the Nash equilibrium relatively easy. The easiest way to determine if a player has a strictly dominant strategy is to find the strategy the player could use that would lead to his or her highest payoff. For Player 1 the highest possible payoff is 18, which occurs when Bottom is used. However, when Player 2 plays Right, Bottom is worse than Middle, and no better than Top, so Bottom is not a strictly dominant strategy. For Player 2 the highest possible payoff is 11, which occurs when Right is used. However, when Player 1 plays Middle, Right is worse than Left, and when Player 1 plays Bottom, Right is worse than either Left of Center.

So neither play has a strictly dominant strategy. But does either player have a strategy that would never be used, or, in the terminology we have, does either player have a strictly dominated strategy? If so, we could "remove" it from the game because (1) the player who has a strictly dominated strategy knows that he will never use it and (2) because all players observe all strategies and payoffs, the other player in the game also knows that a strictly dominated strategy will not be used. In this game, "Top" is strictly dominated by "Middle" for Player 1 ($8 > 7$, $10 > 6$, $6 > 4$). So we know (more importantly, Players 1 and 2 know) that Player 1 will never choose top and so we can eliminate it from the game. The game is now:

		Player 2		
		Left	Center	Right
Player 1	Middle	$8, 8$	$10, 4$	$6, 7$
	Bottom	$18, 7$	$11, 9$	$4, 6$

Once "Top" is removed the strategy "Right" is strictly dominated by the strategy "Left" for Player 2. Note that this was not the case when "Top" was still considered by Player 1. So now we eliminate "Right" and get:

		Player 2	
		Left	Center
Player 1	Middle	$8, 8$	$10, 4$
	Bottom	$18, 7$	$11, 9$

We can now see that "Bottom" is a strictly dominant strategy for Player 1, and Player 1 will use this strategy. Also, we can say that "Middle" is strictly dominated by "Bottom" and eliminate "Middle". This leaves:

		Player 2	
		Left	Center
Player 1	Bottom	$18, 7$	$11, 9$

It is quite obvious that Player 2 will choose "Center" because $9 > 7$. Thus, the solution, or Nash equilibrium, to this game is Player 2 chooses "Bottom" and Player 1 chooses "Center". This method of eliminating strictly dominated strategies is known as **iterated elimination of strictly dominated strategies (IEDS)**.

3 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, $\sigma_i : S_i \to [0, 1]$, 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 $b_i : S_{-i} \to S_i$ 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 correspondence (or function in some specific games we will discuss later).

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.

3.1 Finding PSNE – best responses

Before moving on to a game with mixed strategies, let us now reconsider a game we have already seen:

		Player 2		
		Left	Center	Right
	Top	$7,4$	$6,3$	$4,11$
Player 1	Middle	$8,8$	$10,4$	$6,7$
	Bottom	$18,7$	$11,9$	$4,6$

We just defined a Nash equilibrium in pure strategies (or a "pure strategy Nash equilibrium – PSNE) as a set of strategies such that all players are playing best responses to each other. It is quite easy in these 2-player games with a small number of strategies to determine a player's best response correspondence. Simply fix a strategy for one player and determine what the other player would choose if he knew the other player was using the fixed strategy. If Player 1 chooses Top, Player 2's best response is to choose Right because 11 is the highest payoff. If Player 1 chooses Middle, Player 2 would choose Left, and if Player 1 chooses Bottom Player 2 would choose Center. Thus, Player 2's best response correspondence is $b_2(Top, Middle, Bottom) = \{Right, Left, Center\}$. Player 1's best response correspondence is $b_1(Left, Center, Right) = \{Bottom, Bottom, Middle\}$. We already know from IEDS that $(Bottom, Center)$ is the Nash equilibrium, but finding the best responses also bears this out. From these best responses we also see that Top is *never* a best response by Player 1, and when we used IEDS we eliminated Top.

An easier method of determining the Nash equilibrium using best responses is to simply "mark" the payoffs that correspond to the best responses, as I have done in the matrix below:

		Player 2		
		Left	Center	Right
	Top	$7,4$	$6,3$	$4,\boxed{11}$
Player 1	Middle	$8,\boxed{8}$	$10,4$	$\boxed{6},7$
	Bottom	$\boxed{18},7$	$\boxed{11},\boxed{9}$	$4,6$

Again, notice that no payoffs for Player 1 are marked for the strategy Top, and that the only outcome cell with both payoffs marked is $(Bottom, Center)$. That is the same Nash equilibrium we found using IEDS. While it is unlikely that you will be building formal models of this type (or, if you are, you will likely not be publishing them in top journals), the goal is to understand the intuition for these simple games.

3.2 Finding MSNE

Let's discuss 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* $s_i' \notin S_i^+$

Proof. It is the first proof that we are discussing in this class. Recall that "if and only if" means that the conditions are necessary and sufficient.

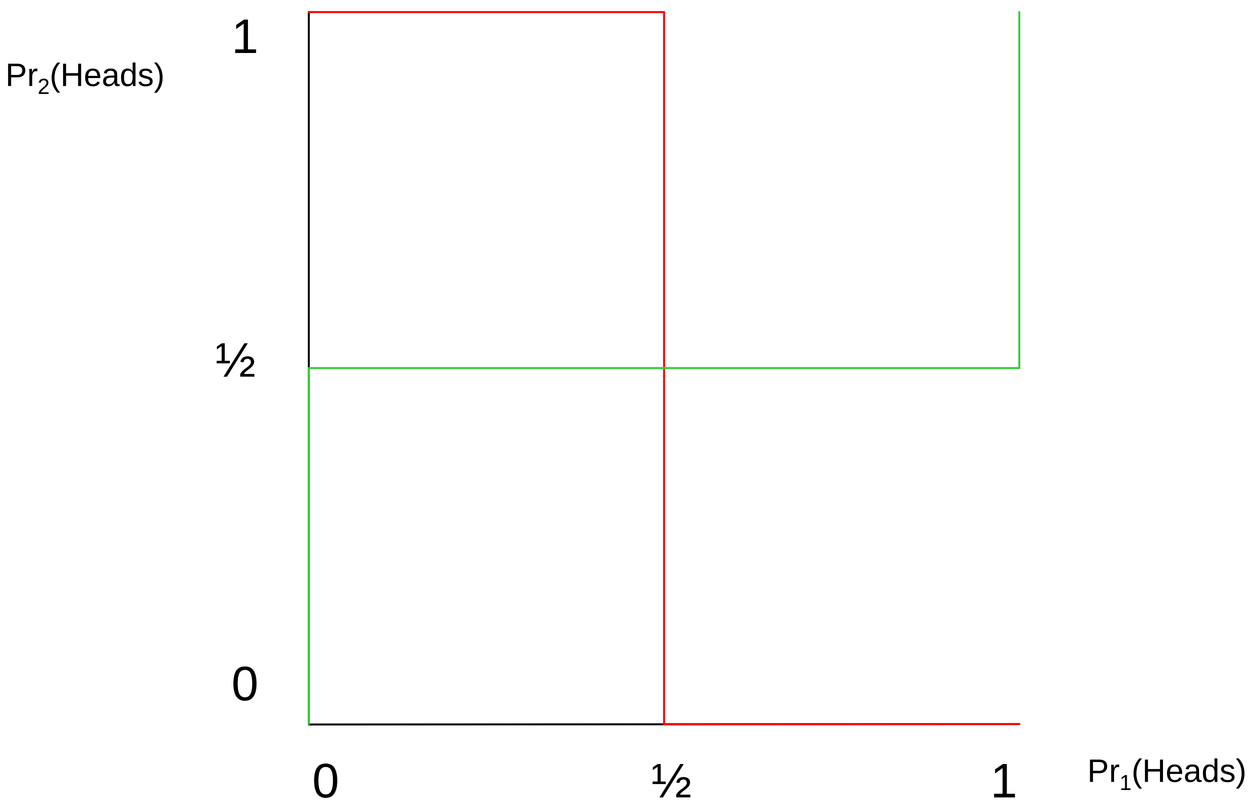

Figure 2: Best response correspondences for the Matching Pennies game.

For necessity, if 1 or 2 do not hold for some player i, then there must be some $s_i \in S_i^+$ and $s_i' \in S_i$ such that player i could switch from s_i to s_i' and receive a strictly higher payoff.

For sufficiency, we use proof by contradiction. Suppose that 1 and 2 hold but σ is not a Nash equilibrium. Because σ is not a Nash equilibrium, that means some player i has some σ_i' such that $u_i(\sigma_i', \sigma_{-i}) > u_i(\sigma, \sigma_{-i})$. If that is true, then player i must have some $s_i' \in S_i^+$ played in σ_i' such that $u_i(s_i', \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i})$. But $u_i(\sigma_i, \sigma_{-i}) = u_i(s_i, \sigma_{-i})$ for $s_i \in S_i^+$, which contradicts 1 and 2. ■

In words, this proposition states that all players must be indifferent between all of their pure strategies over which they assign positive probability, and that the utility from those pure strategies to which a zero probability is assigned must be equal to or less than the utility from a pure strategy played with positive probability. 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:

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

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, the boxing match (Boxing) or the opera (Opera). 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 Boxing to meeting at Opera and Player 2 prefers meeting at Opera to meeting at Boxing. The matrix representation is below:

		Player 2	
		Boxing	Opera
Player 1	Boxing	$3,2$	$0,0$
	Opera	$0,0$	$2,3$

In this game it is easy to see that there are two pure strategy Nash Equilibria. One is Player 1 chooses Boxing and Player 2 chooses Boxing and the other is Player 1 chooses Opera and Player 2 chooses Opera. However, there is also a mixed strategy Nash Equilibrium to this game, where Player 1 chooses Boxing 60% of the time and Opera 40% of the time and Player 2 chooses Boxing 40% of the time and Opera 60% of the time. The graph of the best response correspondences is here, with the green line being the best response of P1 and the red line being the best response of P2.:

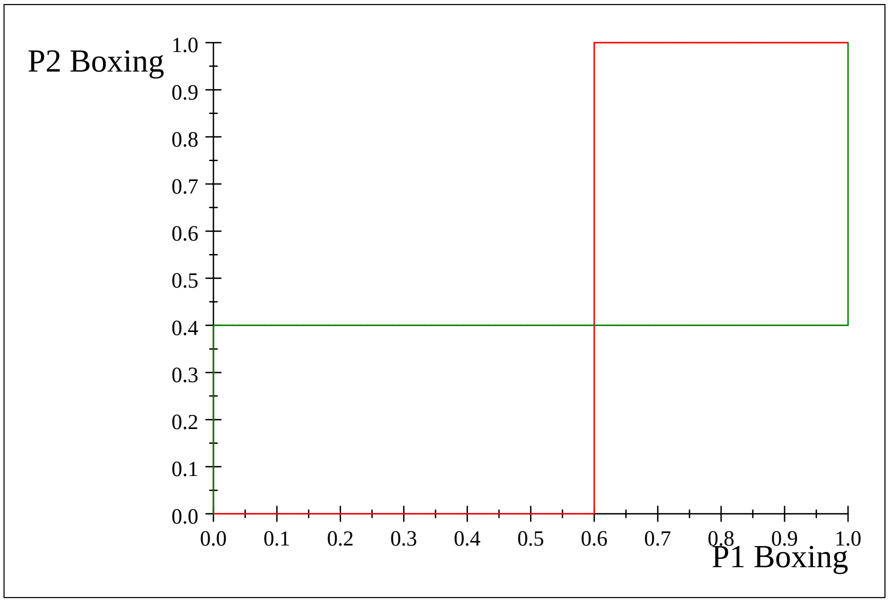

Notice that there are three intersection points that correspond to the three equilibria (2 PSNE, 1 MSNE).

3.3 Existence

In standard consumer and producer optimization problems you may have determined properties under which an equilibrium exists and under which it is unique. 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).*

Note the underlying assumptions in the proposition. First, the number of players is finite (I). Second, the *pure* strategy spaces of the players are all finite, though they need not have the same amount of potential strategies (one player might have 3 strategies, another player 103). Third, we need to allow for mixed strategies. We have already seen that Matching Pennies, a 2-player game where each player has 2 pure strategies, does not have a pure strategy Nash equilibrium. So, if we allow for mixed strategies and the finite number of players has a finite number of pure strategies to choose from then we are guaranteed to have at least one Nash equilibrium. Figure 2 provides an illustration of this concept. Nash (1950) provides the proof in about one page.

Perhaps the most difficult thing to understand in the basic game is the continuity of the payoff functions. Recall that the argument of $u_i(\cdot)$ is the strategy profile, $(\sigma_1, ..., \sigma_I)$. Small changes in Player 1's strategy should not cause large changes in Player 1's payoffs given other players' strategies. Consider Matching Pennies, and fix Player 2's strategy at Heads. A "small change" in Player 1's strategy would be moving from choosing Heads 50% of the time to choosing Heads 50.01% of the time (perhaps that is too large of a change - we can make it 50.0000001% of the time). By moving from 50% Heads, 50% Tails to 50.01% Heads, 49.99% Tails, Player 1's expected payoff increases from 0 to 0.0002. Smaller changes in strategy would lead to smaller changes in payoffs.

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 (note the subtle difference in the statement of this proposition with the statement of the prior one). Footnote 5 on page 253 of MWG is helpful here. We will come back to this second proposition when we discuss games with a continuum of strategies.

These results hinge on fixed-point theorems. The particular fixed-point theorem (Brouwer, Kakutani, etc.) used depends on the assumptions one makes about the strategy sets and payoff functions. 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. We will not go through the formal proofs in class though they are in Appendix 8.A in MWG or section 1.3 of FT (or Nash's 1950 paper). The first proposition is essentially a special case of the second. An important point to remember is that a Nash equilibrium will exist for any game that meets these assumptions. However, there are games which do not meet these assumptions in which a NE exists, it is just that you cannot use these propositions to show existence.