Finding Nash Equilibria

Brett Devine

Games

Games are formalized *metaphors*, or abstractions of real world strategic situations. Games specify who the players in the game (or just number of players) are that will be engaging in interdependent decision making according to the rules of the game. Rules specify the timing at which players act (take their turn to "make a move"), the permitted actions for each player on each of their turns and also what players know or believe at each of their turns. The players act with purpose as the way they jointly play the game leads to outcomes. The game specifies, for each player, preferences over all possible outcomes of the game. A player's preferences are typically represented by a "payoff" to each player in every possible outcome of the game.

The Extensive Form (Game Trees)

Lets first consider a super simple game, called Matching Pennies, involving two players, Alice and Bob. Each player has a penny and they choose whether to lay the penny on a table Heads side up, or Tails side up. If both Alice and Bob play the same side up, Alice wins. However, if they place different sides of the penny facing up then Bob wins. This outcomes are represented by a payoff of 1 for winning and a -1 for losing. The initial node of the game tree has Alice making a move. After Alice makes her choice, Bob will make his choice. The game pictured below is a game of perfect information.

In our description of games, we said that games describe the players, when the players move, what actions players have when it is their turn to move and *what a player believes or what information the player has when they get to move*. In this first game Alice's single decision node is the only one in her information set. She knows that she is the first person to move in the game and the Bob will move second. Bob has two decision nodes representing that his turn comes after Alice chooses either Heads or Tails. Why two decision nodes? Because Bob gets to make his choice of Heads or Tails after two possible contingencies - Alice chooses Heads and Alice chooses Tails. As modeled, Bob gets to observe Alice's choice before making his own. Hence, when Bob makes his choice he knows whether he is moving at the decision node following Alice's choice of Heads or the decision node following Alice's choice of Tails. For this reason we will say that Bob has two information sets, each containing one decision node.

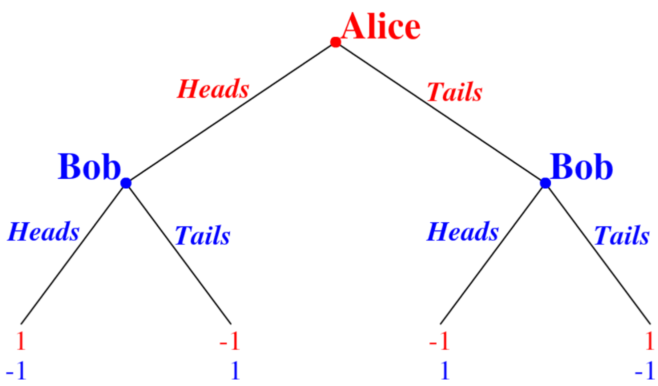

A singleton is a set containing only a single element. For example were we to denote Alice's decision node as n_A and Bob's decision nodes as n_H and n_T (H for Alice choosing Heads and T for Alice choosing Tails) then Alice has a single information set $\{n_A\}$ and Bob has two information sets $\{n_H\}, \{n_T\}$. When all information sets, for all players in the game are singletons, then we say the game is one of perfect information.

How will this game play out? If Alice chooses Heads, Bob easily wins by playing Tails. If Alice plays Tails, Bob easily wins again by playing Heads. In fact, this isn't much of a game since Bob will always win with ease.

The game is far more interesting when neither player knows their opponent's move before playing their own. In other words, the players should move *simultaneously*. From an information perspective whether the players count, "one, two, three, go!" or whether they each play their move covered with their hand and only reveal their choices after both have made their choice the strategic aspect is the same. The chronological timing of moves is less important than the knowledge of what move another player made prior to decision making.

We represent the lack of knowledge a player has about another player's move by combining the relevant decision nodes into the same information set. When we transform the Matching Pennies game into a simultaneous game, the game tree is drawn as below.

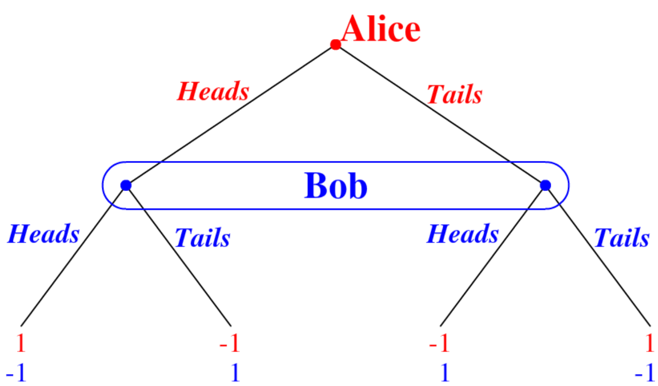

From the new game tree, we see that Alice still has a single information set with one decision node. Now, however, instead of two information sets Bob has only one information set which contains two decision nodes. Bob knows when it is his turn, because he knows which information set he is called to move at. However, because both decision nodes are contained in the same information set Bob cannot be sure which decision node he is actually moving at – meaning he is ignorant of Alice's move.

Pure Strategies

Now we have described that games specify the players, the timing of turns and what players both can do on their turn and what they know (or believe) about other's actions before choosing their own actions as well as the payoffs associated with game outcomes. We now consider the different ways a player can play the game, or in other words, what strategies a player could use to play through the entire game. A strategy provides a player with complete and unambiguous instructions of how to behave should any unexpected contingencies arise during play.

Definition 1. *A **strategy** is a complete, contingent plan for how to play through a game.*

Complete, because we *specify a single action at every information set* that a player has in the game and thus specifies what action the player takes whenever it is their turn. Contingent, because we *specify actions even at information sets that are never reached* due to the actions of any of the players.

For the first version of the matching pennies game, Alice has one information set. There are two possible actions available to Alice at this single information set. Therefore, Alice has exactly two strategies: *Heads* and *Tails*. We call the set of all possible strategies for a player in a game the player's **strategy set**. In other words, Alice's strategy set is $S_A = \{Heads, Tails\}$.

Now lets find Bob's strategies. Bob has two information sets $\{n_H\}$ and $\{n_T\}$ which represents separate contingencies so we need to specify an action for Bob should Alice play *Heads* and also an action should Alice choose *Tails*. Without loss of generality we will order Bob's information sets as $\{n_H\}/\{n_T\}$. Since Bob has 2 information sets, each with 2 available actions, he has $2 \times 2 = 4$ possible strategies. Bob's strategy set is then $S_B = \{Heads/Heads,\ Heads/Tails,\ Tails/Heads,\ Tails/Tails\}$.

Let us now return to the second version of the Matching pennies game with incomplete information. As before, Alice has only a single information set with two possible actions and her strategy set is unchanged $S_A = \{Heads,\ Tails\}$. Bob on the other hand now has only 1 information set with two possible actions. His strategy set is now reduced to $S_B = \{Heads,\ Tails\}$.

Lets look at another slightly more complicated generic game between Alice and Bob.

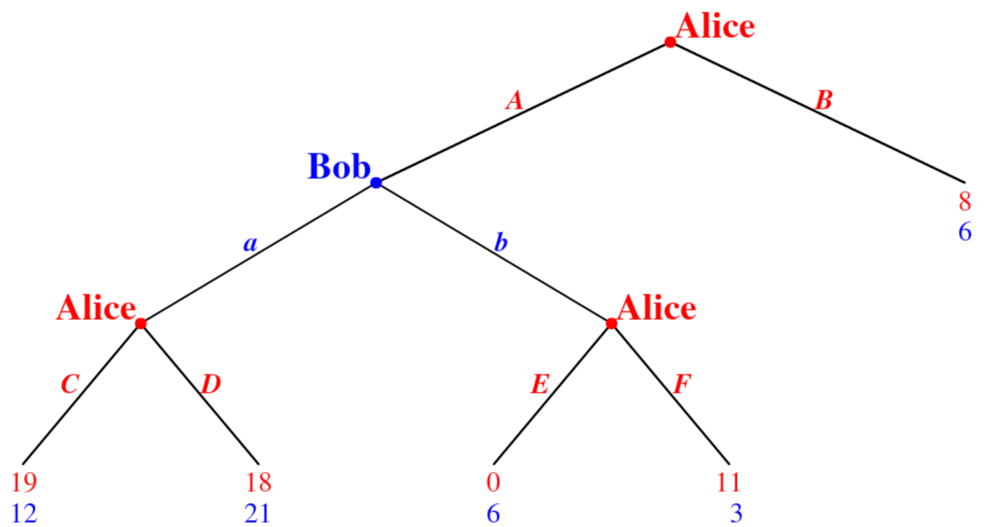

In the above game Alice moves first, then Bob and then Alice again. Alice can choose to end the game on her first turn by selecting action B. If she plays A then Bob will have a turn to select an action. Lets find the strategy sets for both players, starting with Alice.

Alice has three information sets, all of which are singletons. At each of these information sets Alice has 2 possible actions meaning she will have $2 \times 2 \times 2 = 8$ distinct strategies in this game. Proceeding through the game tree from top to bottom and left to right, her strategy set is:

$$S_A = \{A/C/E,\ A/C/F,\ A/D/E,\ A/D/F,\\ B/C/E,\ B/C/F,\ B/D/E,\ B/D/F\}$$

Suppose that Bob decides that if he ever gets to move (meaning Alice doesn't choose B) then he is going to play a. Then Alice will never have an opportunity to play actions E and F. However, a strategy must be a complete contingent plan and specify an action at every information set, even if it is never reached. We do this "just in case" Bob ends up choosing b on his turn. Under that *contingency* we need to make sure that Alice has a plan, or instructions about how to act.

Bob has only a single information set in this game with 2 actions meaning he has 2 possible strategies and a strategy set $S_B = \{a,\ b\}$.

Now consider what happens when Bob can hide his choice from Alice.

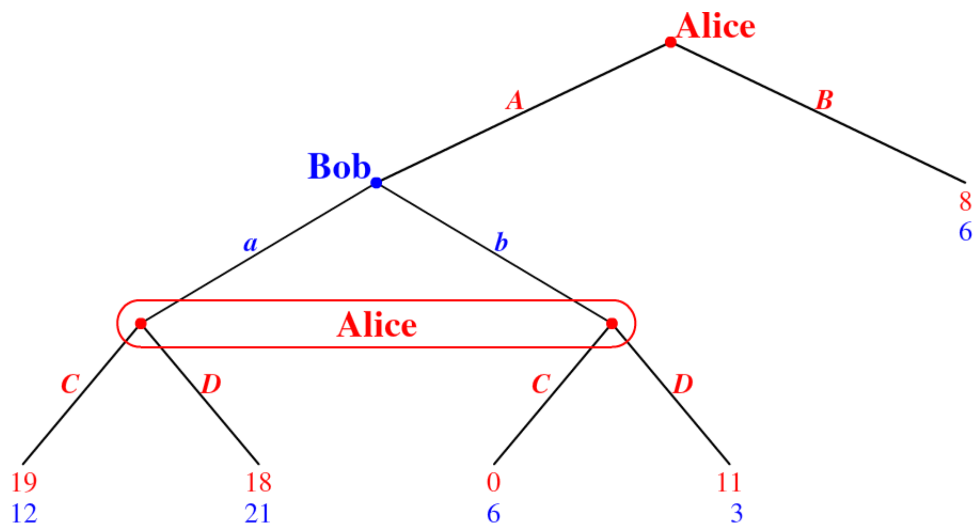

SPECIAL NOTE: We removed Alice's actions *E* **and** *F*. By combining Alice's last two decision nodes into a single information set we are trying to signify that Alice doesn't know whether Bob chose action *a* or *b* before she makes her choice of action. However, there is a problem. Alice's available actions at each node are completely different. If Alice is ever presented with a choice of *E* and *F*, then she knows that Bob chose action *b*. Similarly, if presented with *C* and *D*, she nows Bob chose action *a*.

Rule 1. *An important rule of imperfect information: A player MUST have the SAME set of actions available at every decision node contained in the same information set.*

In this incomplete information version, Alice has 2 information sets each with 2 actions and therefore has $2 \times 2 = 4$ distinct strategies.

$$S_A = \{A/C,\ A/D,\ B/C,\ B/D\}$$

Wait, why are we specifying any further action for Alice if her first move in the game is *B*? She ends the game if she chooses *B*, right? Strategies are complete, contingent plans which specify an action at every information set, even if that information set is never reached. One important reason, is that Bob will need to react to what she would do if she gets to move at her second information set. Why does that matter if she ends the game though? Well, because Bob's reaction to her strategies specification of her actions at her second information set might affect her actions at her first information set based on her expectations of Bob's subsequent actions – i.e., whether she "wants" to choose *A* or *B*. We will discuss strategy selection later. Finally, Bob's strategy set is $S_B = \{a,\ b\}$.

Lets keep going with some more complicated games.

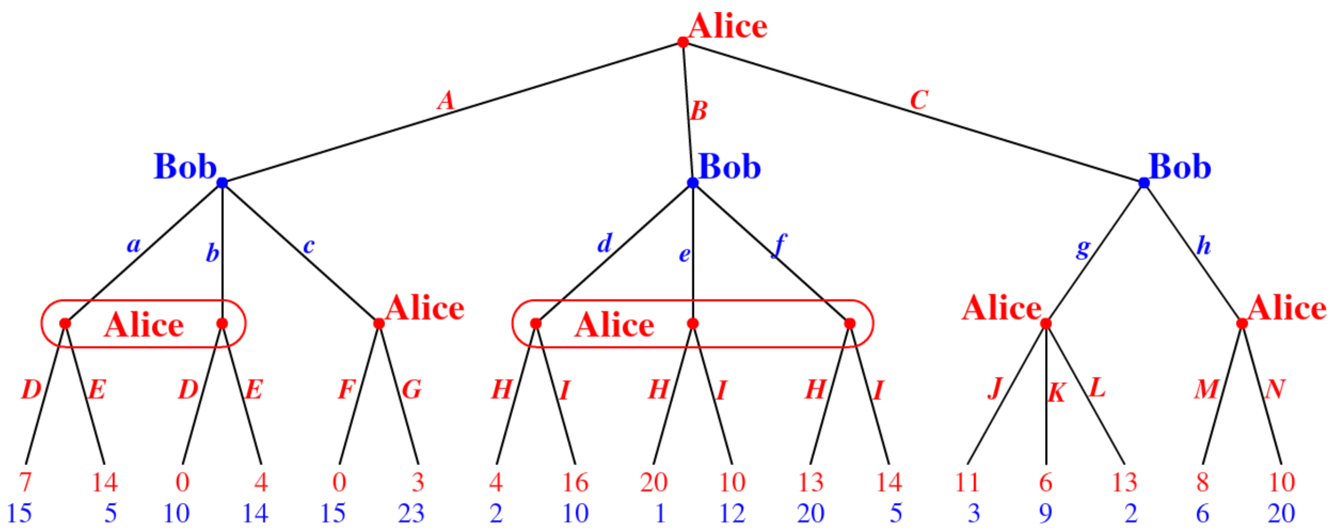

Now Alice and Bob are playing a more complicated game of incomplete information. How many information sets does Alice have? 6. Moving down the game tree and flowing from left to right, how many actions are available to Alice in each of her 6 information sets?

1. 3 - $\{A, B, C\}$
2. 2 - $\{D, E\}$
3. 2 - $\{F, G\}$
4. 2 - $\{H, I\}$
5. 3 - $\{J, K, L\}$
6. 2 - $\{M, N\}$

How many distinct strategies does Alice have? She has $3 \times 2 \times 2 \times 2 \times 3 \times 2 = 3^2 \times 2^4 = 9 \times 16 = 144$ distinct strategies! While it would be quite a burden to list all of these strategies out by hand it will have the general appearance of

$$S_A = \begin{Bmatrix} A/D/F/H/J/M, & B/D/F/H/J/M, & C/D/F/H/J/M, \\ A/E/F/H/J/M, & B/E/F/H/J/M, & C/E/F/H/J/M, \\ A/D/G/H/J/M, & B/D/G/H/J/M, & C/D/G/H/J/M, \\ \vdots & \vdots & \vdots \end{Bmatrix}$$

How many information sets does Bob have? 3. He as 3 actions at his first information set, 3 at his second information set, and 2 actions at his third information set. Hence, Bob has $3 \times 3 \times 2 = 9 \times 2 = 18$ distinct strategies.

$$S_B = \begin{Bmatrix} a/d/g, & b/d/g, & c/d/g, \\ a/e/g, & b/e/g, & c/e/g, \\ a/f/g, & b/f/g, & c/f/g, \\ a/d/h, & b/d/h, & c/d/h, \\ a/e/h, & b/e/h, & c/e/h, \\ a/f/h, & b/f/h, & c/f/h \end{Bmatrix}$$

Intuitively, strategies represent the ways a player can play the game in which they're involved. In the the complete information version of matching pennies, Alice has only 2 different ways to play the game, while Bob has 4 different ways to play. In the incomplete information (simultaneous) version each player has only two ways to play the game.

In the next complete information game, Alice had 8 different strategies, or 8 different ways to play while Bob had only 2. Finally the super complicated game gave Alice a whopping 144 different ways to play while Bob was limited to a mere 18 ways to play.

Continuous Strategy Spaces

We will now apply the same principles we have been using to a more complicated case in which players in the game have continuous action spaces. Consider the simultaneous Bertrand pricing game. There are $N = 2$ firms (players) and both of the firms will simultaneously choose a price to sell their products at. The firms sell an identical product and prices can be chosen to be as small as 0 and as large as 100, or anything in between. To represent this we say a firm's price p_i must be chosen from the interval of real numbers $[0, 100]$. Market demand is given by $100 - p$. If firm 1's price p_1 is less than firm 2's price p_2 (i.e., $p_1 < p_2$) then firm 1 will sell $100 - p_1$ units at price p_1 and firm 2 will sell 0 units. If instead $p_2 < p_1$ then firm 1 will sell 0 units while firm 2 will get the whole market selling $100 - p_2$ units at price p_2. Finally, if the two firms choose the same price, $p_1 = p_2$ then they split the market and each firm sells $(1/2)(100 - p)$ units at the common price $p = p_1 = p_2$.

While strategy sets represent the different ways a single player can approach the game, how many ways can the game be played by all players collectively? Once a single strategy is selected for each player in the game, each player has a complete contingent approach to playing the game. The combined strategies will dictate a complete playing of the game and lead to an outcome and payoffs for the players.

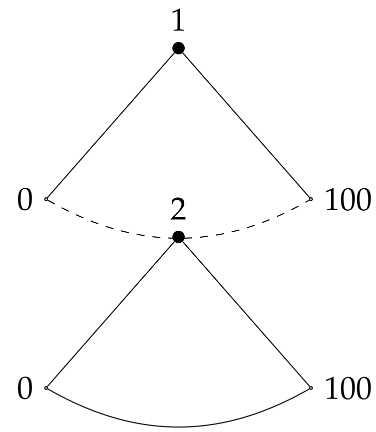

At firm 1's decision node, it has actions $[0, 100]$ available. How many branches (actions) coming out of the decision node are possible? There are an uncountably infinite number. However, we obviously cannot draw that many branches coming out of the first node. Instead we have two branches, one for the lower bound 0 and another for the upper bound 100. We connect these two branches together with a curved line representing all the possible "nodes" the uncountably infinite number of choices of firm 1 could lead to. After any choice $p_1 \in [0, 100]$, firm 2 will take their turn to choose a price, $p_2 \in [0, 100]$. With the previous simple discrete games we modeled earlier, every action by the first player leads to a decision node for player 2. In this pricing game, how many decision nodes does firm 2 have? The answer is an uncountably infinite number - a decision node for every possible choice by firm 1 $p_1 \in [0, 100]$. However, firm 2 does not observe firm 1's choice before picking their own price p_2. Therefore, firm 2 would not know which decision node they are at when making a decision and all of the possible decision nodes are contained in a single information set for firm 2. The curved dashed line connecting the lower and upper bounds of firm 1's choices represent that the infinite number of decision nodes are contained in the same information set.

What if we allowed firm 2 to observe firm 1's choice prior to picking their own price (a game of perfect information)? We could redraw the game as:

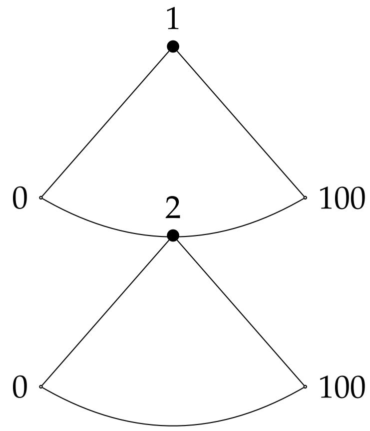

where now the solid curved line connecting the lower and upper bounds of firm 1's action set represents the infinite number of decision nodes that firm 2 can distinguish between. Each of the decision nodes for player 2 are contained in their own set. Firm 1 still has a single information set, but how many information sets does firm 2 have now? – an uncountably infinite number of information sets.

How do we characterize strategies for these players? Since firm 1 has a single information set, a strategy for player 1 is simply a number p_1 in the interval $[0, 100]$. Forming a strategy for player 2 is a different story. We still must specify an action $p_2 \in [0, 100]$ for every information set of player 2, which are characterized by $p_1 \in [0, 100]$. So a strategy for player 2 will specify a price p_2 for every price p_1 in the interval $[0, 100]$. What this means is that a strategy for player 2 will be a function that spits out a specified price $p_2 \in [0, 100]$ whenever we put a price p_1 into it. In other words if $p_2 = f(p_1)$ specifies a price for every $p_1 \in [0, 100]$ then the function $f(p_1)$ represents a strategy for player 2.

I know this discussion of games with continuous strategies sounds complicated, but when we start deriving best response functions using calculus it will make more sense. Essentially there are an infinite number of possible strategies for firm 2, just as there are for firm 1. Best response functions will allow us to identify the "best" strategy for the firms.

Pure Strategy Profiles

When players each have multiple strategies (ways they can approach the game) there will be multiple ways the game will be collectively played by all players. How many different ways could the game be played by everyone together? We call these combinations of strategies across the game's players **strategy profiles**. The question is important, because one of the main goals of game theory is to try and predict how the game will be played by everyone together (which strategy profile we expect to observe) given assumptions about things like rationality, common knowledge of rationality and other factors.

Definition 2. *For a game with $N \geq 1$ players, a* ***strategy profile*** *is a list of strategies, $(s_1, s_2, \ldots, s_N)$ where $s_i \in S_i$ for $i = 1, 2, \ldots, N$ containing exactly one strategy for every player in the game taken from their respective strategy sets.*

Game 1

Lets start off simple with a game of $N = 2$ players where players have rather simple strategy sets of $S_1 = \{a, b\}$ and $S_2 = \{x, y\}$.

The **strategy profiles** of the game are *lists* of strategies, *one for each player*, that take the form (s_1, s_2) where s_1 is some strategy for player 1 from his strategy set $S_1 = \{a, b\}$ and s_2 is some strategy for player 2 from her strategy set $S_2 = \{x, y\}$.

How many of these strategy profiles (pairs of strategies) are there? The first slot can take 2 values (a and b) and the second value can also take on 2 values (x and y). Therefore we have $2 \times 2 = 4$ possible combinations representing all possible strategy profiles for this simple game which are all listed below. The total number of strategy profiles represents to the number of different ways the game can be played by all players together.

(a, x)	(a, y)
(b, x)	(b, y)

Each strategy profile represents a way the players can jointly play the game and leads to an outcome that provides a payoff to each player $u_i(s_1, s_2)$. Therefore each strategy profile (s_1, s_2) results in a *pair* of payoffs $(u_1, u_2) = (u_1(s_1, s_2), u_2(s_1, s_2))$. If we replace the 4 strategy profiles above with the pairs of payoffs that they each lead to, then we get the following structure, called the **strategic form** of the game.

	x	y
a	(a,x)	(a,y)
b	(b,x)	(b,y)

$\longrightarrow$

	x	y
a	(3,3)	(0,5)
b	(5,0)	(1,1)

Game 2

Now consider a slightly more involved game with $N = 2$ players having strategy sets $S_1 = \{T, M, B\}$ and $S_2 = \{L, R\}$. Again, a strategy profile will be a pair (2-tuple) (s_1, s_2) where s_1 is a value from $\{T, M, B\}$ and s_2 is a value from $\{L, R\}$. How many possible strategy profiles are there? Well the first value of the pair can take 3 values while the second value can take 2 values. Hence there are $3 \times 2 = 6$ possible strategy profiles.

(T,L)	(T,R)
(M,L)	(M,R)
(B,L)	(B,R)

Each of these strategy profiles will result in a complete play of the game and a pair of payoffs (a payoff for each player) of the form $(u_1(s_1, s_2), u_2(s_1, s_2))$. The set of strategy profiles can again form a matrix (strategic form game).

	L	R
T	(T,L)	(T,R)
M	(M,L)	(M,R)
B	(B,L)	(B,R)

$\longrightarrow$

	L	R
T	(3,2)	(-1,2)
M	(5,5)	(0,0)
B	(0,0)	(1,1)

Game 3

To come full circle consider the game of market entry posed below. There exists a market with an incumbent monopolist and a potential entrant is considering market entry. The incumbent will need to decide whether to accommodate the firm's entrance (conceding market share) or fight the firm's entrance at possibly great cost.

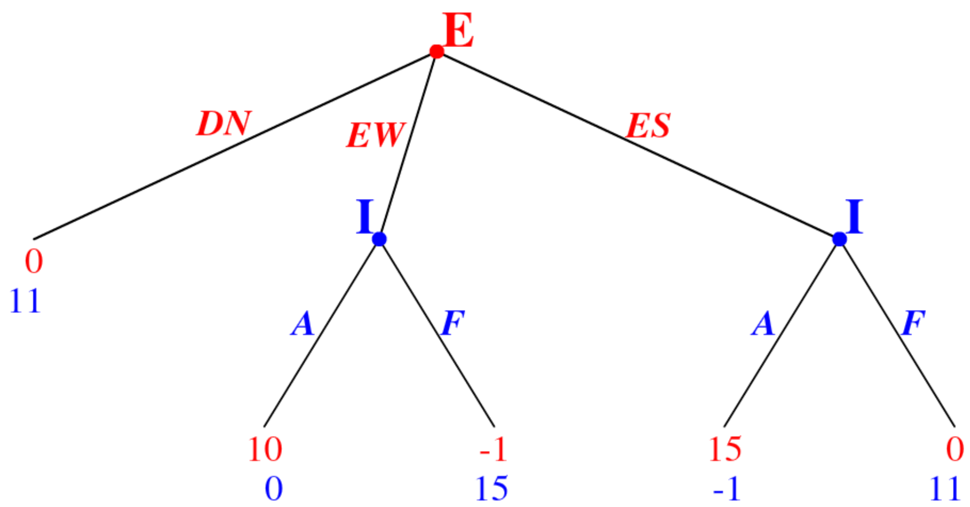

The potential entrant moves first at a single information set containing a single decision node. The entrant's available actions are DN = Do not enter, EW = Enter Weak, and ES = Enter Strong. Since this is the entrant's only information set these available actions also represent the entrant's strategy set $S_E = \{DN, EW, ES\}$.

The incumbent has two information sets, each containing a single decision node and has actions A = Accommodate and F = Fight. The incumbent has 2 information sets with 2 actions at each and therefore $2 \times 2 = 4$ distinct pure strategies resulting in the strategy set $S_I = \{A/A,\ A/F,\ F/A,\ F/F\}$. Given these strategy sets, there are $3 \times 4 = 12$ strategy profiles and the set of all possible strategy profiles is

(DN,A/A)	(DN,A/F)	(DN,F/A)	(DN,F/F)
(EW,A/A)	(EW,A/F)	(EW,F/A)	(EW,F/F)
(ES,A/A)	(ES,A/F)	(ES,F/A)	(ES,F/F)

Using the payoffs from the extensive form game, the strategy sets and the possible strategy profiles we can construct the strategic form of the game.

		Incumbent			
		A/A	A/F	F/A	F/F
	DN	$0,11$	$0,11$	$0,11$	$0,11$
Entrant	EW	$10,0$	$10,0$	$-1,15$	$-1,15$
	ES	$15,-1$	$0,11$	$15,-1$	$0,11$

Game 4

Lets get a little more complex for this next game. Let the number of players be $N = 3$ and the strategy sets are

$$S_1 = \{a, b, c, d, e\}$$
$$S_2 = \{m, n, o, p\}$$
$$S_3 = \{x, y\}$$

Since we have $N = 3$ players, each strategy profile will be a list (3-tuple) of strategies (s_1, s_2, s_3) where s_1 is a strategy from player 1's strategy set S_1, s_2 is a strategy from player 2's strategy set S_2 and s_3 is a strategy from player 3's strategy set S_3. How many possible strategy profiles are there in this game? The first item s_1 can take on 5 values, the second item s_2 can take on 4 values while the third item s_3 can take on 2 values. Hence, there are $5 \times 4 \times 2 = 40$ distinct strategy profiles representing 40 different ways the 3 players can collectively play the game. The set of strategy profiles for this game are listed below.

(a,m,x)	(a,n,x)	(a,o,x)	(a,p,x)	(a,m,y)	(a,n,y)	(a,o,y)	(a,p,y)
(b,m,x)	(b,n,x)	(b,o,x)	(b,p,x)	(b,m,y)	(b,n,y)	(b,o,y)	(b,p,y)
(c,m,x)	(c,n,x)	(c,o,x)	(c,p,x)	(c,m,y)	(c,n,y)	(c,o,y)	(c,p,y)
(d,m,x)	(d,n,x)	(d,o,x)	(d,p,x)	(d,m,y)	(d,n,y)	(d,o,y)	(d,p,y)
(e,m,x)	(e,n,x)	(e,o,x)	(e,p,x)	(e,m,y)	(e,n,y)	(e,o,y)	(e,p,y)

How can than this game be represented in the strategic form? 2 player games can be conveniently represented by a matrix, but this is a three player game. To represent this three player game we will use two matrices, each corresponding to a strategy choice by player 3. To make this transition clear look at the set of strategy profiles for this game listed previously. There are two 5 by 4 "blocks" of strategy profiles. Look closely at the last strategy in each profile (the one corresponding to player 3's strategy). In the left most block all the strategy profiles involve player 3 playing x. In the right most block all the strategy profiles involve player 3 playing y. In both blocks, the rows index player 1's strategies and the columns index player 2's strategies. We have not specified payoffs for this "game" but when creating the strategic form each cell of the two matrices will correspond to unique strategy profiles leading to some (unspecified) payoffs for each player u_1, u_2, u_3.

Player 2 (Player 3 = x)

		m	n	o	p
	a	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3
	b	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3
Player 1	c	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3
	d	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3
	e	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3

Player 2 (Player 3 = y)

		m	n	o	p
	a	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3
	b	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3
Player 1	c	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3
	d	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3
	e	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3	u_1, u_2, u_3

Game 5

Consider a game with $N \geq 2$ players where each player has an identical strategy set $S_i = \{mac, windows\}$ for $i = 1, 2, \ldots, N$. A strategy profile for this game will be a list (n-tuple) specifying a single strategy for each player. Each item of the n items in the strategy profile can each take on 2 values. Hence, the total number of strategy profiles is $2 \times 2 \times \cdots \times 2 = 2^N$. For clarity, if $N = 2$ there are $2^2 = 4$ strategy profiles. If $N = 3$ there are $2^3 = 8$ strategy profiles and $N = 5$ would result in $2^5 = 32$ strategy profiles describing how the game can be played.

Game 6

Now consider a game with N players where M of the N players have the strategy set $S_i = \{a, b, c\}$ and L of the players have the strategy set $S_j = \{x, y, z\}$ and the remaining $N - L - M$ players have the strategy set $S_k = \{0, 1, 2, 3, 4, 5, 6, 7\}$. How many possible strategy profiles are there? For M players their strategy can take 3 values, for L players their strategies can take on 3 values and the remaining $N - M - L$ players have strategies that take on 8 possible values. Hence, we have $3^M \times 3^L \times 8^{N-M-L}$ possible strategy profiles.

Game 7

What about pure strategy profiles when the players have continuous strategies as in the Bertrand pricing game discussed previously? In the imperfect simultaneous version of the game each firm has one information set and a strategy was simply a number p in the interval $[0, 100]$. A pure strategy profile in this case is simply a pair of numbers (p_1, p_2) where both p_1 and p_2 are in the interval $[0, 10]$.

What about the sequential perfect information version? Firm 1's strategy was just a number p_1 in the interval $[0, 10]$. Firm 2's strategies where functions $f(p_1)$ of firm 1's choice. These functions take a number in the interval $[0, 10]$ and return a number in $[0, 100]$ representing firm 2's price. What do the strategy profiles look like then? Well there are two players, so all strategy profiles will be a list (pair) of two items, the first a number p_1 for firm 1 and the second a function $f(p_1)$ for firm 2. In other words we will have strategy profiles $\{(p_1, f(p_1))\}$ for all $p_1 \in [0, 100]$ and for every possible function $f : [0, 100] \rightarrow [0, 100]$. How many different strategy profiles are there? An uncountably infinite number.

Symmetric and Asymmetric Strategy Profiles

We now consider a special case of games in which all players have the same strategy sets. Return to the simultaneous matching pennies game between Alice and Bob. Both players have the same strategy set $S = \{Heads, Tails\}$. The strategy profiles were simply

(Heads,Heads)	(Heads, Tails)
(Tails,Heads)	(Tails, Tails)

Definition 3. *If all players in a game have identical strategy sets then we say a strategy profile is a* ***symmetric strategy profile*** *if all players are playing the same strategy.*

Lets apply this definition to the simultaneous matching pennies game. The two strategy profiles $(Heads, Heads)$ and $(Tails, Tails)$ have every player in the game playing the same strategy and are therefore symmetric strategy profiles.

Definition 4. *If all players in a game have identical strategy sets then we say a strategy profile is an* ***asymmetric strategy profile*** *if at least one player plays a different strategy than the other players.*

Applying this definition to matching pennies, the asymmetric strategy profiles are clearly $(Heads, Tails)$ and $(Tails, Heads)$.

Lets consider a slightly larger game. Imagine a game with $N = 4$ players who each have an identical strategy set $S_i = \{T, M, B\}$ for $i = 1, 2, 3, 4$. The strategy profiles for this game will be 4-tuples containing one strategy for each player. How many possible strategy profiles are there? each slot in the strategy profile has 3 possible values and there are 4 players so we have $3 \times 3 \times 3 \times 3 = 3^4 = 81$ possible strategy profiles. How many of these 81 strategy profiles are symmetric strategy profiles? Well, by definition all players must be playing the same strategy in any symmetric strategy profile. By this we know that there are only 3 possible symmetric strategy profiles.

(T, T, T, T)
(M, M, M, M)
(B, B, B, B)

All of the remaining $81 - 3 = 78$ strategy profiles have at least one player (maybe more) playing a different strategy than the other players which by definition makes them asymmetric strategy profiles.

Operating System Game Consider again the operating system game in which $N \geq 2$ consumers choose between two different computer platforms, Mac or Windows. This game has N players and each has the same strategy set $S_i = \{Mac, Windows\}$ for $i = 1, 2, \ldots, N$. We already know that strategy profiles in this game will be a list of N strategies, one for each player (N-tuples). We also know that there are 2^N possible strategy profiles. But how many symmetric strategy profiles are there? There are only 2 symmetric strategy profiles in this game. One in which every player chooses *Mac* and one in which every player chooses *Windows*. All other $2^N - 2$ strategy profiles have at least one consumer choosing a computer platform that is different than other players and are therefore asymmetric strategy profiles.

Bertrand Pricing Game The imperfect information Bertrand pricing game with continuous strategy spaces discussed previously had $N = 2$ players and the players had symmetric strategy sets, $S_i = [0, 100]$ for $i = 1, 2$. We have already discussed how there are an uncountably infinite number of possible strategy profiles. But how many symmetric strategy profiles are there and what do they look like? Well, a symmetric strategy profile in this game will be a pair (p_1, p_2) such that $p_1 = p_2$. This means that for every possible number p in $[0, 100]$ there is a symmetric strategy profile (p, p). Hence, there are an uncountably infinite number of symmetric strategy profile. Any profile (p'_1, p'_2) where $p'_1 \neq p'_2$ is an asymmetric strategy profiles.

A Key Point of Game Theory

Up to this point we have had a long discussion of describing games, deriving each player's strategy sets and finally the possible strategy profiles which describe all the ways players could jointly play the game. We have also worked on listing and counting all the possible strategy profiles. In some cases there are a lot of strategy profiles.

One of the key points, or goals of game theory is to generate reasonable predictions about which of the possible many strategy profiles in a game we might expect to observe – i.e., predict how players will collectively end up playing the game.

We investigate different "solution methods" which reduce the set of reasonable strategy profiles down to a small enough number that we can get useful predictions. For example, under the assumption of rationality, we know rational players will never play strictly dominated strategies. Consider a game with $N = 3$ players with strategy sets $S_1 = \{a, b, c, d\}$, $S_2 = \{x, y, z\}$ and $S_2 = \{m, n, o\}$. We can calculate that there are $4 \times 3 \times 3 = 36$ possible strategy profiles for this game. But suppose that we can determine that d is a strictly dominated strategy for player 1 and that z is a strictly dominated strategy for player 2. Then if players 1 and 2 are both rational we know that any strategy profile in which player 1 is playing d or in which player 2 is playing z is not a reasonable prediction and we remove those strategy profiles. This would leave us with only $3 \times 2 \times 3 = 18$ remaining strategy profiles – cutting the number in half!

The most important "solution concept" is that of Nash equilibria and its various refinements and extensions that we will cover later in the course (subgame perfection, bayesian, perfect bayesian, etc.) **Nash equilibrium is a definition that applies to strategy profiles**. We seek to find Nash equilibrium strategy profiles.

Nash Equilibria

Lets again start with a simple 2 player game between Alice and Bob who have found themselves arrested for committing a crime. The detectives have placed them in separate interrogation rooms. The detectives don't have enough evidence to convict them both for all their crimes in a trial and they need to get a confession.

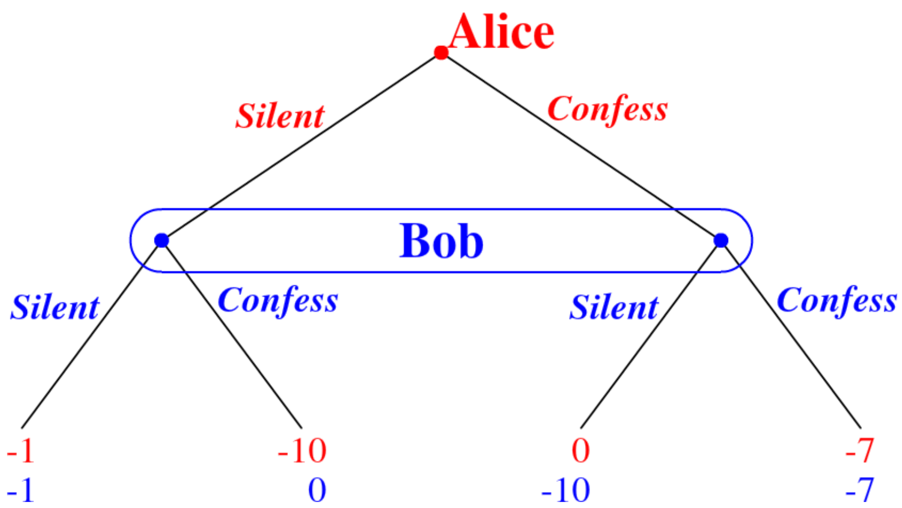

Both Alice and Bob have an identical strategy set $S_i = \{Silent, Confess\}$. They will not know the other's choice of strategy before making their own choice. However, the detectives truthfully inform them of the payoffs of their consequences. If both Alice and Bob choose to stay silent (strategy profile $(Silent, Silent)$) then they both serve only 1 year in prison – a payoff of -1. If both choose to confess, then they will both go to prison for 7 years – a payoff of -7. However, the detectives make them a deal that if they are the only one to confess, then they will be set free today and their partner in crime will go to prison for 10 years.

The strategy profiles of this game are

(Silent, Silent)	(Silent, Confess)
(Confess, Silent)	(Confess, Confess)

Unilateral Deviation from Strategy Profile

We now discuss what it means for a player to *unilaterally deviate* from a given strategy profile of a game. Suppose we propose a strategy profile $(Silent, Silent)$ and consider what Alice can do to change the strategy profile all by herself. There are only two strategy profiles in the game in which Bob plays $Silent$, and those are $(Silent, Silent)$ and $(Confess, Silent)$. For Alice a unilateral deviation from $(Silent, Silent)$ is a switch to the strategy profile $(Confess, Silent)$. Similarly for Bob, a unilateral deviation involves changing strategy profiles from $(Silent, Silent)$ to $(Silent, Confess)$.

Given any strategy profile for a game, a unilateral deviation involves changing exactly one strategy in the given strategy profile. Why would a player ever want to deviate from the given strategy profile to another in which only their own strategy is different? – because they believe they can gain a strictly higher payoff by doing so.

Definition 5. *(2-Player Version) For a $2-$Player game, a strategy profile $s^* = (s_1^*, s_2^*)$ is a* ***Nash equilibrium*** *if*

$$u_1(s_1^*, s_2^*) \geq u_1(s, s_2^*) \text{ for all } s \in S_1$$
$$u_2(s_1^*, s_2^*) \geq u_2(s_1^*, s) \text{ for all } s \in S_2$$

Please take some time to digest what the above definition is really saying. If a specific strategy profile in a game has the property that NO PLAYER in the game CAN DO STRICTLY BETTER by being the only one to change their strategy from the one specified for them in the strategy profile (unilateral deviation), then that strategy profile is a Nash equilibrium. This strategy profile (a way for the players to collectively play the game) is considered stable because none of the players wants to be the only one to change their strategy (unilaterally deviate). Since no player wants to be the only one to change strategies, none of them will change strategies to try and do better in the game - so the strategy profile is not prone to changes.

Lets apply the definition now to Alice and Bob's prisoner's dilemma. Lets propose the following strategy profile to check against the definition, $(Silent, Silent)$. First, we consider Alice's possible unilateral deviations $(Confess, Silent)$ and we compare her payoffs. If she sticks with the proposed strategy profile Alice will receive a payoff of $u_{Alice}(Silent, Silent) = -1$. If she instead unilaterally deviates to the profile $(Confess, Silent)$ then she will receive a payoff of $u_{Alice}(Confess, Silent) = 0$. This is a problem since we now have

$$u_{Alice}(Silent, Silent) < u_{Alice}(Confess, Silent)$$

Because of this, the strategy profile $(Silent, Silent)$ fails to satisfy the definition of a Nash equilibrium. Alice isn't the only one interested in deviating either, since Bob would prefer a unilateral deviation to the profile $(Silent, Confess)$.

Okay, so Alice wanted to move to the profile $(Confess, Silent)$, is this a Nash Equilibrium? First lets consider Alice once again by comparing potential payoffs

$$u_{Alice}(Confess, Silent) = 0 > u_{Alice}(Silent, Silent)$$

which satisfies the first condition that Alice cannot do strictly better by unilaterally deviating, so she won't. But what about Bob? If Bob sticks with the proposed strategy profile his payoff is $u_{Bob}(Confess, Silent) = -10$ while if he unilaterally deviates he receives payoff $u_{Bob}(Confess, Confess) = -7$. Since

$$u_{Bob}(Confess, Confess) > u_{Bob}(Confess, Silent)$$

Bob can do strictly better by deviating to the profile $(Confess, Confess)$. Therefore the strategy profile $(Confess, Silent)$ also cannot be a Nash equilibrium.

Well what about the profile $(Confess, Confess)$? A unilateral deviation by Alice to $(Silent, Confess)$ will strictly reduce her payoff, so she prefers to stay at $(Confess, Confess)$. A unilateral deviation by Bob to $(Confess, Silent)$ will strictly reduce his payoff so he also prefers to stay at $(Confess, Confess)$. Therefore, the symmetric strategy profile $(Confess, Confess)$ is a Nash equilibrium. In fact, the profile $(Confess, Confess)$ is the only Nash equilibrium in the game.

Matching Pennies ... Again Lets apply the definition of Nash equilibrium to the simultaneous matching pennies game discussed several times previously.

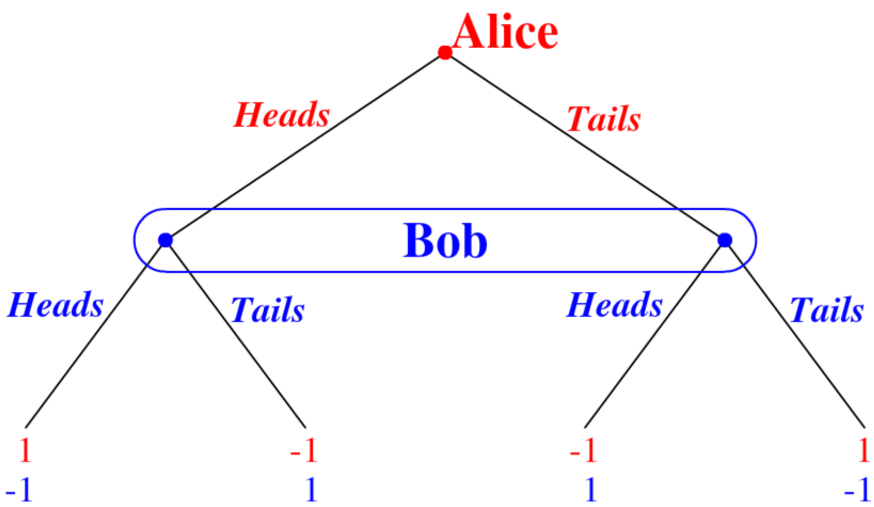

Lets begin with the symmetric strategy profiles. Could $(Heads, Heads)$ be a Nash equilibrium? If Alice sticks with this profile she receives a payoff of $u_{Alice}(Heads, Heads) = 1$. If she were to unilaterally deviate to $(Tails, Heads)$, then her payoff would be $u_{Alice}(Tails, Heads) = -1$ which is worse than staying at $(Heads, Heads)$. Hence, Alice does not want to be the only one to switch strategies. Now we check for Bob. If Bob stays at $(Heads, Heads)$ then his payoff is $u_{Bob}(Heads, Heads) = -1$ whereas if he unilaterally deviates to $(Heads, Tails)$ is payoff increases to $u_{Bob}(Heads, Tails) = 1$. Because Bob can do strictly better by unilaterally deviating from $(Heads, Heads)$ we know $(Heads, Heads)$ is not a Nash equilibrium.

We now test $(Tails, Tails)$. Bob's payoff is $u_{Bob}(Tails, Tails) = -1$ if he stays at $(Tails, Tails)$ but he receives a payoff of $u_{Bob}(Tails, Heads) = 1 > -1$ by unilaterally deviating to $(Tails, Heads)$. Therefore $(Tails, Tails)$ cannot be a Nash equilibrium.

What about $(Heads, Tails)$? Under this profile Alice receives $u_{Alice}(Heads, Tails) = -1$ but she can do strictly better by unilaterally deviating to $(Tails, Tails)$ and getting payoff $u_{Alice}(Tails, Tails) = 1$ so $(Heads, Tails)$ cannot be a Nash equilibrium.

There is only one strategy profile remaining, $(Tails, Heads)$. Under this profile Alice receives $u_{Alice}(Tails, Heads) = -1$ but if she unilaterally deviates to $(Heads, Heads)$ then Alice can do strictly better and receive $u_{Alice}(Heads, Heads) = 1$. So $(Tails, Heads)$ is also not a Nash equilibrium. But wait, we tested all of the pure strategy profiles, what happened? For some games, there is no pure strategy Nash equilibrium profile. Later we will discuss mixed strategy profiles and in that case there is always at least one mixed strategy Nash equilibrium profile.

Battle of the Sexes Alice and Bob are trying to meet up for a date, but can't communicate about which of two venues to attend. Both have identical strategy sets $S_i = \{Opera, Fight\}$.

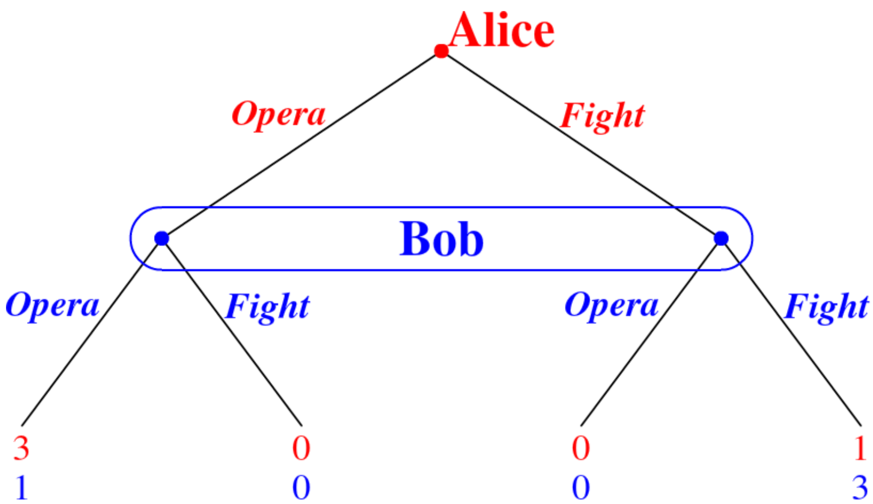

Two players with 2 strategies mean we have $2 \times 2 = 4$ strategy profiles

(Opera, Opera)	(Opera, Fight)
(Fight, Opera)	(Fight, Fight)

There are 2 symmetric strategy profiles $(Opera, Opera)$ and $(Fight, Fight)$ and 2 asymmetric strategy profiles $(Opera, Fight)$ and $(Fight, Opera)$. Are there any pure strategy Nash equilibria? Lets test the asymmetric profiles first – consider $(Opera, Fight)$. Under this strategy profile Alice receives $u_{Alice}(Opera, Fight) = 0$ and we compare this with the payoff she could get by unilaterally deviating to another strategy profile. The only other possible strategy profile with Bob playing *Fight* is $(Fight, Fight)$ so we compare Alice's payoff under $(Opera, Fight)$ to her payoff under $(Fight, Fight)$ and we see

$$u_{Alice}(Opera, Fight) = 0 < 1 = u_{Alice}(Opera, Opera)$$

which violates a requirement of a Nash equilibrium and therefore Alice benefits from unilaterally deviating. Therefore the profile $(Opera, Fight)$ is not a Nash equilibrium.

What about $(Fight, Opera)$? If we look for unilateral deviations from this strategy profile that make a player strictly better off we will find that both Bob and Alice have such deviations and hence $(Fight, Opera()$ violates the definition of a Nash equilibria.

Lets look at the symmetric profiles, starting with $(Opera, Opera)$. Alice's payoff is $u_{Alice}(Opera, Opera) = 3$ and if she unilaterally deviates she receives $u_{Alice}(Fight, Opera) =$

0 which makes her strictly worse off. Hence we have the condition

$$u_{Alice}(Opera, Opera) \geq u_{Alice}(s, Opera) \text{ for all } s \in S_A$$

What about Bob? Bob's payoff is $u_{Bob}(Opera, Opera) = 1$ and his payoff from unilaterally deviating is $u_{Bob}(Opera, Fight) = 0$. Thus a profitable deviation for Bob does not exist from $(Opera, Opera)$ and we have the condition

$$u_{Bob}(Opera, Opera) \geq u_{Bob}(Opera, s) \text{ for all } s \in S_B$$

and therefore $(Opera, Opera)$ satisfies the definition of a Nash equilibrium.

What about the remaining symmetric profile $(Fight, Fight)$? Alice's payoff is 1 and she would get 0 if she deviates. Bob's payoff from $(Fight, Fight)$ is 3 and he would get 0 if she unilaterally deviates. Since no one strictly benefits from deviating, no one will deviate and $(Fight, Fight)$ is a Nash equilibrium.

The prisoner's dilemma game, the matching pennies game and the battle of the sexes game all had 4 possible strategy profiles. However, the matching pennies game had zero pure strategy Nash equilibria, the prisoner's dilemma game has 1 pure strategy Nash equilibrium and the battle of the sexes game has two pure strategy Nash equilibria.

Important Tip Look at the definition of a Nash equilibrium strategy profile closer and realize that it requires that ALL players in the game DO NOT HAVE a unilateral deviation that makes them STRICTLY better off. Given any strategy profile, if you can find JUST ONE player that has a unilateral deviation that makes that player strictly better off then you can conclude the given strategy profile is NOT a Nash equilibrium.

Lets move on to a three player game.

Definition 6. *(3-Player Version) For a 3-Player game, a strategy profile $s^* = (s_1^*, s_2^*, s_3^*)$ is a **Nash equilibrium** if*

$$\begin{aligned} u_1(s_1^*, s_2^*, s_3^*) &\geq u_1(s_1, s_2^*, s_3^*) \textit{ for all } s_1 \in S_1 \\ u_2(s_1^*, s_2^*, s_3^*) &\geq u_2(s_1^*, s_2, s_3^*) \textit{ for all } s_2 \in S_2 \\ u_3(s_1^*, s_2^*, s_3^*) &\geq u_3(s_1^*, s_2^*, s_3) \textit{ for all } s_3 \in S_3 \end{aligned}$$

Where S_1 is player 1's strategy set, S_2 is player 2's strategy set, and S_3 is player 3's strategy set.

Again take some time to make sure you understand every little part of the above definition. The definition applies to strategy profiles that are 3-tuples (ordered list of 3 strategies) such as (s_1, s_2, s_3). It says that for a Nash equilibrium strategy profile, we shouldn't be able to find a single player with a unilateral deviation that makes that player strictly better off. If we can, then the profile is not a Nash equilibrium by definition. If we can't, then that profile is a Nash equilibrium.

Lets practice with the following 3 player game below.

Player 3: I

Player 1 \ Player 2	x	y	z
a	2,1,2	0,0,2	1,2,3
b	0,3,1	2,2,4	3,1,0
c	1,1,1	3,2,1	2,2,2

Player 3: II

Player 1 \ Player 2	x	y	z
a	2,1,3	1,0,3	1,0,4
b	1,2,1	3,3,3	1,1,1
c	1,2,1	1,0,0	2,1,2

The game has $N = 3$ players and each player's strategy sets are as follows.

$$S_1 = \{a, b, c\}$$
$$S_2 = \{x, y, z\}$$
$$S_3 = \{I, II\}$$

How many strategy profiles are there? $3 \times 3 \times 2 = 18$. The 18 strategy profiles are listed below.

(a,x,I)	(a,y,I)	(a,z,I)	(a,x,II)	(a,y,II)	(a,z,II)
(b,x,I)	(b,y,I)	(b,z,I)	(b,x,II)	(b,y,II)	(b,z,II)
(c,x,I)	(c,y,I)	(c,z,I)	(c,x,II)	(c,y,II)	(c,z,II)

Notice in the layout of the strategy profiles how the rows correspond to strategies of player 1, the columns to strategies of player 2 and the different matrices (groupings of 9 strategies) correspond to player 3's two strategy choices.

Lets pick a random profile and test it, say (a, x, I). Player 1's payoff is $u_1(a, x, I) = 2$. The possible unilateral deviations for player 1 are to (b, x, I) and (c, x, I). Neither of these profiles provide a strictly higher payoff to player 1, so far so good. Player 2's payoff is $u_2(a, x, I) = 1$ and player 2's possible unilateral deviations are to profiles (a, y, I) and (a, z, I) which yield payoffs of 0 and 2 respectively. Since $u_2(a, x, I) = 1 < 2 = u_2(a, z, I)$ player 2 can do strictly better by deviating and therefore (a, x, I) is not a Nash equilibrium.

Lets try another one. I propose the strategy profile (c, y, I). Starting with player 1, the possible unilateral deviations are to (a, y, I) and (b, y, I). Player 1 is currently receiving a payoff of 3 from the profile (c, y, I) and can receive 2 and 0 from respectively deviating to (b, y, I) and (a, y, I). Since player 1 cannot do strictly better by deviating, we cannot yet reject the proposed strategy profile as a Nash equilibrium. Under (c, y, I) player 2 is receiving a payoff of 2, but would receive payoffs 1 and 2 by unilaterally deviating to (c, x, I) or (c, z, I) respectively. Since player 2 cannot do *strictly* better by unilaterally deviating, player 2 will stay at (c, y, I) and we proceed to check player 3. Player 3 currently receives a payoff of 1. Unilaterally deviating to (c, y, II) would yield a payoff of 0 to player 3 and hence player 3 will not want to deviate. Since no player can do strictly better by unilaterally deviating we know the conditions in the 3-Player definition of Nash equilibrium are satisfied and (c, y, I) is a Nash equilibrium strategy profile.

You might be thinking this is a lot of work to go through the 18 different strategy profiles and check each one against the definition of Nash equilibrium. Is there a better way? Sometimes. The method of "testing" a strategy profile against the definition is critically important to be comfortable and skilled with as it can be useful in the widest number of games. However, the **best response** (or best-reply) method can sometimes be faster and sometimes the only practical way of finding Nash equilibrium strategy profiles.

The best-response method involves finding the strategy that maximizes a player's private payoff given the player's correct beliefs about the strategies the other players will select. In terms of strategy profiles, we specify the subset of strategy profiles that have the other player's strategies fixed. From the 3 player game above, we could specify the all the strategy profiles in which players 2 and 3 play (z, I) which are the profiles

(a,z,I) (b,z,I) (c,z,I)

Player 1's strategy set is $S_1 = \{a, b, c\}$ and so we can determine which strategy gives player 1 the highest payoff given beliefs that players 2 and three are playing strategies z and I. Given these beliefs, b is the strategy that maximizes player 1's payoff, making it his best response. We can highlight player 1's payoff in the game matrix. Next suppose player 1 believes that the other players are going to select y and I. Then the relevant strategy profiles are $(a, y, I), (b, y, I), (c, y, I)$. Given these beliefs, strategy c maximizes player 1's payoff and we highlight the payoff 3 in in the payoff cell corresponding to profile (c, y, I).

Player 3: *I*

Player 2

Player 1		*x*	*y*	*z*
	a	2,1,2	0,0,2	1,2,3
	b	0,3,1	2,2,4	**3**,1,0
	c	1,1,1	**3**,2,1	2,2,2

Player 3: *II*

Player 2

Player 1		*x*	*y*	*z*
	a	2,1,3	1,0,3	1,0,4
	b	1,2,1	3,3,3	1,1,1
	c	1,2,1	1,0,0	2,1,2

Lets continue for every possible strategy profile and highlight player 1's best responses. Remember, that in this game player 1 is choosing among the rows, while player 2's strategy specifies the column and player 3's strategy specifies the matrix. After highlights the game matrices appear as follows.

Player 3: *I*

Player 2

Player 1		*x*	*y*	*z*
	a	**2**,1,2	0,0,2	1,2,3
	b	0,3,1	2,2,4	**3**,1,0
	c	1,1,1	**3**,2,1	2,2,2

Player 3: *II*

Player 2

Player 1		*x*	*y*	*z*
	a	**2**,1,3	1,0,3	1,0,4
	b	1,2,1	**3**,3,3	1,1,1
	c	1,2,1	1,0,0	**2**,1,2

We now find the best responses for player 2. Similar to player 1, we consider all the possible combinations of strategies for players 1 and 3 (specifying the row and matrix) while player 2 will choose the strategy (column) that maximizes his payoff. At each stage we will highlight the payoff in the game matrix corresponding to player 2's optimal choice.

Player 3: *I*

Player 1 \ Player 2	*x*	*y*	*z*
a	2,1,2	0,0,2	1,2,3
b	0,3,1	2,2,4	3,1,0
c	1,1,1	3,2,1	2,2,2

Player 3: *II*

Player 1 \ Player 2	*x*	*y*	*z*
a	2,1,3	1,0,3	1,0,4
b	1,2,1	3,3,3	1,1,1
c	1,2,1	1,0,0	2,1,2

Lets proceed to highlight the best response profiles for player 3. Suppose player 3 believes players 1 and 2 to be playing (a, x) then we look at the payoff matrices and see there are two cells in which players 1 and 2 play (a, x) and they relate to profiles (a, x, I) and (a, x, II). Strategy II gives player 3 a payoff of 3 while I yields a payoff of 2 so player 3's best response is II and we highlight player 3's payoff in the strategy profile (a, x, II). Continuing this process for all the possible strategy profiles we get the final highlighted game matrices below.

Player 3: *I*

Player 1 \ Player 2	*x*	*y*	*z*
a	2,1,2	0,0,2	1,2,3
b	0,3,1	2,2,4	3,1,0
c	1,1,1	3,2,1	2,2,2

Player 3: *II*

Player 1 \ Player 2	*x*	*y*	*z*
a	2,1,3	1,0,3	1,0,4
b	1,2,1	3,3,3	1,1,1
c	1,2,1	1,0,0	2,1,2

A strategy profile for the game that is a best response for each player will satisfy the definition of a Nash equilibrium as each player lacks a unilateral deviation that will yield a strictly higher payoff. There are only two strategy profiles in the above highlighted game in which every player's payoff is highlighted, (c, y, I) and (a, x, II). These represent the pure strategy Nash equilibria for the game.

Another 3 Player Game Lets try another 3 player game with slightly larger strategy spaces for player 3. The three players will engage in a vote in which they each can vote for one of three options, A, B and C. The strategy sets for all three players is $S_i = \{A, B, C\}$. This means there are $3 \times 3 \times 3 = 3^3 = 27$ strategy profiles.

3 votes for *A*

		2		
		A	*B*	*C*
	A	2,0,0	2,0,0	2,0,0
1	*B*	2,0,0	1,2,1	2,0,0
	C	2,0,0	2,0,0	0,1,2

3 votes for *B*

		2		
		A	*B*	*C*
	A	2,0,0	1,2,1	1,2,1
1	*B*	1,2,1	1,2,1	1,2,1
	C	1,2,1	1,2,1	0,1,2

3 votes for *C*

		2		
		A	*B*	*C*
	A	2,0,0	0,1,2	0,1,2
1	*B*	0,1,2	1,2,1	0,1,2
	C	0,1,2	0,1,2	0,1,2

We begin by highlighting the best responses for player 1 given the strategies of the other players. Remember again that player 2's strategies specify a column and player 3's strategy specifies a matrix. The highlighted game is below.

3 votes for *A*

1 \ 2	A	B	C
A	2,0,0	2,0,0	2,0,0
B	2,0,0	1,2,1	2,0,0
C	2,0,0	2,0,0	0,1,2

3 votes for *B*

1 \ 2	A	B	C
A	2,0,0	1,2,1	1,2,1
B	1,2,1	1,2,1	1,2,1
C	1,2,1	1,2,1	0,1,2

3 votes for *C*

1 \ 2	A	B	C
A	2,0,0	0,1,2	0,1,2
B	0,1,2	1,2,1	0,1,2
C	0,1,2	0,1,2	0,1,2

Next we will highlight the best response profiles for player 2 given the other's strategies. Player 1's strategies specify a row and player 3's strategy specifies a matrix. Player 2 will then choose the strategy (column) that maximizes his payoff.

3 votes for *A*

2

1		*A*	*B*	*C*
	A	2,0,0	2,0,0	2,0,0
1	*B*	2,0,0	1,2,1	2,0,0
	C	2,0,0	2,0,0	0,1,2

3 votes for *B*

2

1		*A*	*B*	*C*
	A	2,0,0	1,2,1	1,2,1
1	*B*	1,2,1	1,2,1	1,2,1
	C	1,2,1	1,2,1	0,1,2

3 votes for *C*

2

1		*A*	*B*	*C*
	A	2,0,0	0,1,2	0,1,2
1	*B*	0,1,2	1,2,1	0,1,2
	C	0,1,2	0,1,2	0,1,2

Finally, we can proceed to find the best responses for player 3 with the recognition that the strategies of players 1 and 2 will specify a specific cell in each matrix and we will compare player 3's payoffs between identical cells in the three different matrices representing player 3's strategy choices $S_3 = \{A, B, C\}$.

3 votes for A

		2		
		A	B	C
	A	2,0,0	2,0,0	2,0,0
1	B	2,0,0	1,2,1	2,0,0
	C	2,0,0	2,0,0	0,1,2

3 votes for B

		2		
		A	B	C
	A	2,0,0	1,2,1	1,2,1
1	B	1,2,1	1,2,1	1,2,1
	C	1,2,1	1,2,1	0,1,2

3 votes for C

		2		
		A	B	C
	A	2,0,0	0,1,2	0,1,2
1	B	0,1,2	1,2,1	0,1,2
	C	0,1,2	0,1,2	0,1,2

Now that we have identified each player's best responses, we can identify which strategy profiles are best responses for all players in the game. The pure strategy Nash equilibrium are

$$(A,A,A), \quad (B,B,B), \quad (B,B,C), \quad (A,C,C), \quad (C,C,C)$$

General N-player Games We now proceed to apply the definition of Nash equilibrium to games involving an arbitrary number of players. For the moment we will continue to deal with discrete strategy spaces for each of the N players.

Definition 7. *For a game with $N \geq 2$ players each having strategy sets S_i for $i = 1,2,\ldots,N$ a strategy profile $(s_1^*, s_2^*, \ldots, s_N^*)$ is a **Nash Equilibrium** if and only if*

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

for all $i = 1,2,\ldots,N$.

Again take some time to make sure you understand exactly what this definition is saying. s_i represents player i's strategy while s_{-i} is the list of $N-1$ strategies being played by the $N-1$ other players. The condition requires that player i does not have any possible *unilateral deviations* from the strategy profile $(s_1^*, s_2^*, \ldots, s_N^*)$ that will give him a *strictly* greater payoff. This condition needs to be true for all of the N players in the game. In the 2-player definition of Nash equilibrium this amounted to 2 conditions - one for each player. In the 3-player definition of Nash equilibrium this amounted to 3 conditions - one for each player. Now, in the more general case of N players, we have N conditions - one for each of the N players.

Important Tip It is important to understand how to violate the definition of a Nash equilibrium. Why? Because of we can show that just 1 of the N required conditions fails to hold, then we can determine that the particular strategy profile we were testing is not a Nash equilibrium. Again, all we have to do is find one, of the N players, say player j, that under the strategy profile $(s_1^*, \ldots, s_j^*, \ldots, s_N^*)$ has a unilateral deviation to another of his strategies $s_j \in S_j$ that gives him a strictly higher payoff. In other words given the proposed strategy profile $(s_1^*, \ldots, s_j^*, s_N^*)$ and the unilateral deviation profile $(s_1^*, \ldots, s_j, \ldots, s_N^*)$ we have

$$u_j(s_j^*, s_{-j}^*) < u_j(s_j, s_{-j}^*)$$

Now lets apply this general N-player definition of Nash equilibrium to a symmetric N-player game.

Operating System Game Consider the operating system game in which we have N consumers who each have an identical strategy set $S_i = \{Mac, Win\}$. The payoff each consumer receives from each strategy depends on how many other players choose the same strategy. The payoff functions are

$$u_i(Mac) = 100 + 10m$$
$$u_i(Win) = 30(N - m)$$

where m is the number of players choosing *Mac* and $N - m$ of the players choose *Win*.

First note that this is a simultaneous symmetric game since all players have the same strategy sets and the same payoffs from those strategies. Essentially all the players are N copies of each other. Furthermore this game exhibits network effects since the more people choose the same strategy, the better the payoff from that strategy. These network effects often "tip" the equilibrium outcomes toward the symmetric strategy profiles.

How many strategy profiles are there in this game? Each strategy profile will be a list (an N-tuple) of N strategies, one for each player. Since each player has exactly two strategies there are $2 \times 2 \times \cdots \times 2 = 2^N$ possible strategy profiles. How many of these 2^N strategy profiles are symmetric profiles and how many are asymmetric profiles? There are two symmetric strategy profiles, one where everyone chooses *Mac* and one where everyone chooses *Win*.

$$(Mac_1, Mac_2, \ldots, Mac_N) \qquad (Win_1, Win_2, \ldots, Win_N)$$

The remaining $2^N - 2$ strategy profiles are all asymmetric strategy profiles since they involve at least one player choosing a different strategy than everyone else.

Since 2^N is potentially a lot of strategy profiles to check for Nash equilibrium we need to be clever in how we start looking for them. First, we use the symmetry of the game along with the presence of network effects to surmise that the best place to start looking might be the two symmetric strategy profiles. Hence, lets propose the strategy profile in which everyone chooses *Mac*. Now we test this proposed strategy profile against the definition of a Nash equilibrium. We do so by comparing the payoffs a player gets under this proposed profile to what he or she would get if they deviated to another strategy (in this case *Win*).

If all N players choose *Mac* then each player receives a payoff of $100 + 10N$ (since $m = N$). If a player was to unilaterally deviate to *Win* they would be the only player choosing *Win* and their payoff would be $30 \cdot 1 = 30$. Comparing this we see that for everyone choosing *Mac* to be a Nash equilibrium it must be the case that no player can do strictly better by unilaterally deviating to *Win*. This is the case if and only if

$$100 + 10N \geq 30 \tag{1}$$

Note that even if $N = 1$ the left-hand side is always greater meaning that for any number of players, the symmetric strategy profile in which everyone chooses *Mac* is a Nash equilibrium of the game.

What about the symmetric strategy profile where everyone plays *Win*? Well under this profile each player receives a payoff of $30N$. If any of them unilaterally deviate they will be the only player choosing *Mac* and their payoff will be $100 + 10 \cdot 1 = 110$. In order for everyone playing *Win* to be a Nash equilibrium it must be the case that no player can do strictly better by unilaterally deviating to *Mac*. This condition will be true if and only if

$$\begin{aligned} 30N &\geq 110 \\ N &\geq \frac{110}{30} \\ N &\geq 3.67 \end{aligned}$$

Hence, in order for the network effect to be strong enough that everyone choosing *Win* is a Nash equilibrium, we need there to be at least $N = 4$ players. Only under that condition can both symmetric strategy profiles be Nash equilibrium.

What about the asymmetric strategy profiles in this game? Could any of them be pure strategy Nash equilibria? Lets think about this. Any asymmetric profile will involve $1 \leq m < N$ people choosing *Mac* and the remaining $N - m = w$ players choosing *Win*. If any of the m players think they can do better by becoming the $N - m + 1st$ person to choose *Win* then they will unilaterally deviate – ruling it out as a Nash equilibrium. Similarly, if any of the $N - m$ people choosing *Win* believe they can do strictly better by unilaterally deviating to be the $m + 1st$ person choosing *Mac*, then they will do so ruling it out as a Nash equilibrium. Any possible asymmetric Nash equilibrium will need to satisfy a particular property – the number of players choosing each strategy needs to be such that every player is indifferent between their two strategy choices. If indifferent, then deviation cannot make a player strictly better off. *However*, in this particular game we have discrete number of players and we might not have exact indifference. But we will use that as a guide and then test for the specific number of mac and windows users that satisfy the definition of a Nash equilibrium.

Lets proceed. Given m players playing *Mac* and $N-m$ players playing *Win* indifference between *Mac* and *Win* requires that the following equality of payoffs hold.

$$100 + 10m = 30(N - m)$$

Now we solve the above equation for m in terms of N.

$$\begin{aligned} 100 + 10m &= 30(N - m) \\ 100 + 10m &= 30N - 30m \\ 100 + 10m + 30m &= 30N \\ 40m &= 30N - 100 \\ m &= \frac{30N - 100}{40} \end{aligned}$$

We see from the above condition that the value of m depends on the number of players N. In order for m to be at least 1, we require that

$$\begin{aligned} \frac{30N - 100}{40} &\geq 1 \\ 30N - 100 &\geq 40 \\ 30N &\geq 140 \\ N &\geq \frac{140}{30} = 4.67 \end{aligned}$$

Therefore, IF, we are to find an asymmetric Nash equilibrium we will need there to be at least $N = 5$ players.

Now suppose that there are $N = 10$ players.

$$\frac{30(10) - 100}{40} = 5$$

and so we will look at the strategy profile where 5 players choose *Mac* and the remaining $10 - 5 = 5$ players choose *Win*. Lets check to see if this satisfies the definition of a Nash equilibrium. First lets look at mac users. The 5 mac users are currently each receiving a payoff of $100 + 10(5) = 100 + 50 = 150$. If one of them unilaterally deviates to *Win*, they will be become the 6th windows user which would yield a payoff of $30(6) = 180$. Since they can do strictly better by deviating to *Win* we know this cannot be a Nash equilibrium.

What about $N = 20$ users? Well in that case we expect the payoffs to be equal between the two choices when $m = 12.5$. So we can check two different strategy profiles. One where 12 people choose *Mac* and 8 people choose *Win* and a second where 13 people choose *Mac* and 7 people choose *Win*.

If 12 people choose *Mac* their payoff is $100 + 10(12) = 220$. If one of them switches to *Win* they will become the 9th person on windows which means their payoff would be $30(9) = 270$. Since they can do strictly better by becoming the 9th windows user they will unilaterally deviate and this is not a Nash equilibrium.

If 13 people choose *Mac* their payoff is $100 + 10(13) = 230$. Anyone considering switching would become the 8th windows user and would get a payoff of $30(8) = 240$.

Since 240 > 230 there is a unilateral deviation that makes mac users strictly better off so this also is not a Nash equilibrium. What is going on?

Notice that each additional person to windows increases the payoff by 30, while each additional mac user only increases payoffs by 10. This means that when N gets larger, the fixed payoff of 100 from using a mac matters less and less. As a consequence people can often do better for themselves by switching to windows, which causes the game to "tip" toward the symmetric strategy in which everyone uses windows.

What if we looked at smaller numbers? Suppose that $N = 5$ and that 1 person plays *Mac* and $5 - 1 = 4$ people choose *Win*. Then the single mac user gains a payoff of $100 + 10(1) = 110$. By becoming the 5th windows user he would get a payoff of $30(5) = 150$. So unilaterally deviating still makes someone strictly better off.

Think about what we are seeing. Network effects cause player's payoffs from a choice to increase the greater the number of other players also choosing that same strategy. For this reason, whenever a player considers unilaterally deviating and adding one more person to the other strategy she will increase the payoff of everyone choosing that strategy. This means that when we have a number of players m choosing *Mac* and $N - m$ choosing windows such that the payoffs to both groups are either equal, or close to equal, then unilaterally deviating to one side will increase the payoffs of everyone on that side which breaks the near "indifference" and tips the game toward one of the strategy choices.

So does this mean that there is asymmetric Nash equilibrium then? If there is, we would formally require the following conditions to hold for m players choosing *Mac* and $N - m$ choosing *Win*.

$$100 + 10m \geq 30(N - m + 1)$$
$$30(N - m) \geq 100 + 10(m + 1)$$

Let us solve both conditions for m.

$$\begin{aligned} 100 + 10m &\geq 30N - 30m + 30 \\ 70 + 40m &\geq 30N \\ 40m &\geq 30N - 70 \\ m &\geq \frac{30N - 70}{40} \end{aligned}$$

Now for the second condition

$$\begin{aligned} 30(N - m) &\geq 100 + 10(m + 1) \\ 30N - 30m &\geq 100 + 10m + 10 \\ 30N - 110 &\geq 40m \\ \frac{30N - 110}{40} &\geq m \end{aligned}$$

we now have two conditions on m that are inconsistent. Notice that because the denominators are the same and $30N$ is the same we know

$$\frac{30N - 110}{40} < \frac{30N - 70}{40}$$

There does not exist a number m such that $m \leq (30N - 110)/40$ and is also greater than or equal to $(30N - 70)/40$. For this reason we conclude, that there is no asymmetric Nash equilibrium in the game because there is no number m, where $1 \geq m < N$ that will satisfy the no profitable deviation conditions.

A Beautiful Mind In the movie *A Beautiful Mind* there is a scene in a bar in which four men are discussing a young blonde woman who entered the bar with several brunette friends. For some reason, all four men express their desire to date the blonde woman. In particular, they appear to prefer chasing the blonde woman over her brunette friends. One of the men publicly declares the beginning of the competition between all four men to obtain the privilege of dating the blonde. In fact, one of the men (not Russel Crow) attempts to recall a lesson of Adam Smith, the "father of economics" in which he states that "*in competition, individual ambition serves the common good*." The men then toast to the idea "let the best man win!"

At this point, Russel Crow's character John Nash, notices a problem. If all the men chase the blonde, then they will "block" each other and none of them will succeed in obtaining her affections. After blocking each other, all the men will seek the affections of the blonde woman's brunette friends who will reject their advances out of spite for being second choice. As a consequence, all four men will fail to obtain the affections of any woman. John Nash then proposes that all four men agree to not chase the blonde and instead all seek the affections of the brunette woman (of which there are enough for all four men). His claim is that this is what is best both for each individual and for the group. If they all avoid the blonde they will not block each other and because they made the brunettes their first choice they won't be rejected out of spite.

Lets analyze this scenario and see just which strategy profiles are Nash equilibria.

Each of the 4 player have the same strategy set $S_i = \{blonde, brunette\}$. Let b be the number of players that play *blonde*. Then the payoffs for each player are as follows,

$$u_i(blonde) = \begin{cases} 2 & \text{if } b = 1 \\ 0 & \text{if } b > 1 \end{cases}$$
$$u_i(brunette) = 1$$

If a player is the only one to choose *blonde* he gets a payoff of 2, but if even one other does as well, they block each other and both players get 0. On the other hand, a player can always guarantee himself a payoff of 1 by choosing *brunette*.

Each player in the game has 2 strategies, and the strategy profiles will be a list of 4 strategies (one for each player). Since there are 4 players each with 2 strategies we have $2 \times 4 = 8$ possible strategy profiles. Lets first consider the two symmetric strategy profiles.

$$(blonde, blonde, blonde, blonde) \qquad (brunette, brunette, brunette, brunette)$$

When everyone plays *blonde*, $b = 4 > 1$ and each player receives a payoff of 0. For any player, unilaterally deviating to *brunette* will make them strictly better off with a payoff of 1. Therefore, the symmetric strategy profile where everyone plays blonde is not a Nash equilibrium. When everyone plays *brunette* each player receives a payoff of 1. Each player, when considering *unilateral deviations* could switch to *blonde* and be the only one playing *blonde* and will therefore receive a payoff of $2 > 1$ making him strictly better off. Therefore everyone playing *brunette* also fails to be a Nash equilibrium.

We now consider the asymmetric strategy profiles. Consider any strategy profile in which the number of players choosing *blonde* is $b \geq 2$ can be ruled out as a Nash equilibrium as anyone of those $b \geq 2$ players is strictly better off unilaterally deviating to *brunette*. Also we need there to be at least 1 person playing *blonde* so that no one has an incentive to deviate from *brunette* to *blonde*. All this means that the only strategy profiles that could possibly be Nash equilibrium have exactly 1 player playing *blonde* and 3 players playing *brunette*. How many of these profiles are there? 4.

(blonde, brunette, brunette, brunette)	(brunette, blonde, brunette, brunette)
(brunette, brunette, blonde, brunette)	(brunette, brunette, brunette, blonde)

In each strategy profile one player successfully goes for the *blonde* and would do strictly worse $1 < 2$ by switching to *brunette*. Because $b = 1$, if any of the players choosing *brunette* were to unilaterally deviate they would get 0 instead to 1 making them worse off. We see that there does not exist a player who has a unilateral deviation that makes them strictly better off. So each of the four asymmetric strategy profiles above are pure strategy Nash equilibria of the game.

The above game is a symmetric game because all players had the same strategy set and the same preferences over outcomes. However, in equilibrium, the "identical" players make non-identical choices. Why? This is an example of **congestion**. All players prefer *blonde*, but if too many people play that strategy they crowd, or block each other, making it optimal to choose something different. When congestion is a property of a symmetric game, Nash equilibria tend to be asymmetric strategy profiles.

Solution Finding Tip Identify whether a game is a symmetric game. If it is then try to identify, by analyzing payoffs, whether the game is likely to exhibit network effects and tipping or if it will exhibit congestion. If it has properties of network effects and tipping we might expect the Nash equilibrium to be symmetric strategy profiles. On the other hand, if it has properties of congestion we might expect the Nash equilibrium to be in asymmetric strategy profiles.

Cournot Competition We have $N = 2$ firms who compete in producing quantities of a homogenous product. Each firm faces the inverse demand curve $p(Q) = 100 - Q$ where Q is the aggregate market quantity. Firms have symmetric production costs per unit of $c = 10$. Firm's preferences over outcomes are determined entirely by the profit they receive. Both firms have an identical strategy set $S_i = [0, \infty)$ where they choose their quantity $q_i \in [0, \infty)$. Aggregate market quantity is defined as $Q = q_1 + q_2$. The firm's profit functions are then given as

$$\pi_1(q_1, q_2) = (100 - q_1 - q_2)\, q_1 - 10q_1$$
$$\pi_2(q_1, q_2) = (100 - q_1 - q_2)\, q_2 - 10q_2$$

Each firm wants to maximize its profit subject to its beliefs over the strategy of the other player. Suppose that Firm 1 believes that Firm 2 will produce $q_2 = 20$. Then Firm 1's profit maximization problem is

$$\max_{q_1 \in [0,\infty)} (100 - 20 - q_1)q_1 - 10q_1$$

To maximize this function we note that it is concave (hill-shaped) and therefore the first order condition is sufficient for a maximum. This means all we need to concern ourselves with is taking the 1st derivative with respect to q_1, set it equal to zero and solve for q_1 (which finds the spot on the profit curve where the slope equals zero). The first order condition is

$$80 - 2q_1 - 10 = 0$$
$$q_1^* = \frac{70}{2} = 35$$

Okay great, so if Firm 1 believes that Firm 2 will produce $q_2 = 20$, then Firm 1 will maximize profit by choosing $q_1 = 35$. What if beliefs are that $q_2 = 35$? Then the profit maximization problem is

$$\max_{q_1 \in [0,\infty)} (100 - 35 - q_1)q_1 - 10q_1$$

The associated first order condition is

$$65 - 2q_1 - 10 = 0$$

implying the optimal q_1 is $q_1 = 55/2$ which is less than 35.

A key requirement of Nash equilibrium is not just that players maximize their payoff given their beliefs about the strategies of others, but that those beliefs are *correct*. The problem is that we don't know what Firm 2 will do exactly nor what Firm 1 will do for Firm 2 to decide. To go through all possible beliefs and calculate the optimal strategy would be impossible. So instead, we let Firm 1's beliefs about Firm 2 be represented by the variable q_2 and Firm 2's beliefs about Firm 1 by the variable q_1.

So we rephrase Firm 1's profit maximization problem as

$$\max_{q_1 \in [0,\infty)} (100 - q_2 - q_1)q_1 - 10q_1$$

The first order condition is

$$100 - q_2 - 2q_1 - 10 = 0$$
$$90 - q_2 = 2q_1$$
$$q_1^* = \frac{90 - q_2}{2}$$

What about Firm 2?

$$\max_{q_2 \in [0,\infty)} (100 - q_1 - q_2)q_2 - 10q_2$$

The first order condition for the maximum is

$$
\begin{aligned}
100 - q_1 - 2q_2 - 10 &= 0 \\
q_2^* &= \frac{90 - q_1}{2}
\end{aligned}
$$

Each of the functions we solved for $q_1^*(q_2)$ and $q_2^*(q_1)$ are called **best-response functions**. As we stated we need each player's beliefs to be correct *in equilibrium*. Hence, we require that

$$
\begin{aligned}
q_1^* &= \frac{90 - q_2^*}{2} \\
q_2^* &= \frac{90 - q_1^*}{2}
\end{aligned}
$$

Since we know what q_2^* is, we can substitute that into q_1^*.

$$
\begin{aligned}
q_1^* &= \frac{1}{2}\left(90 - \frac{90 - q_1^*}{2}\right) \\
&= \frac{1}{2}\left(\frac{180}{2} - \frac{90 - q_1^*}{2}\right) \\
&= \frac{1}{2}\left(\frac{180 - 90 + q_1^*}{2}\right) \\
&= \frac{90 + q_1^*}{4} \\
4q_1^* &= 90 + q_1^* \\
3q_1^* &= 90 \\
q_1^* &= \frac{90}{3} = 30
\end{aligned}
$$

Now we plug $q_1^* = 30$ into $q_2^*(q_1)$ to represent correct beliefs for Firm 2 and we get $q_2^* = (1/2)(90 - 30) = (1/2)(60) = 30$.

This means that for the strategy profile $(30, 30)$ each firm maximizes their profit given their beliefs about the strategy of the other firm. If either firm unilaterally deviated to any other quantity they would not be able to do strictly better for themselves (by definition of the maximization problem). So the Nash equilibrium strategy profile is $(q_1^* = 30,\ q_2^* = 30)$. Graphically, we can plot the best response functions and identify the Nash equilibrium as the point of mutual best response.

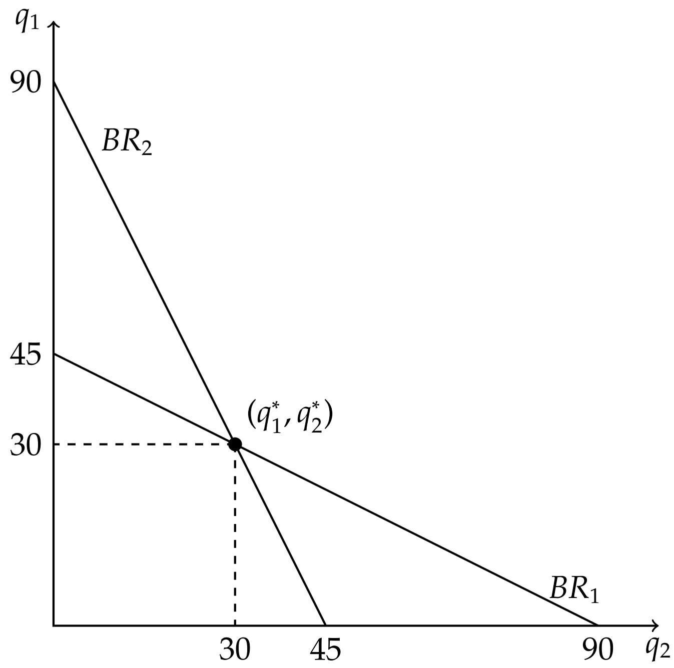

Figure 1: Graph of best response functions in Cournot competition