Game Theory

Mark Dean

Lecture Notes for Fall 2009 Introductory Microeconomics - Brown University

1 Introduction

So far, the agents that we have considered in our economies have not paid much attention to each other. In the case of perfect competition, firms and consumers pretty much completely ignored each other: they both took prices as given, and bought and sold as much as they wanted at those prices. Consumers paid no attention as to how their behavior would affect the behavior of firms, and similarly firms paid no attention to how their actions affected the behavior of consumers. When we discussed monopolistic firms things got a little bit more sophisticated: the firm took into account how the prices they set would affect the choices of the consumer, but that was it: consumers in turn did not in turn consider how their actions might affect firms. One way of saying this is that consumers were not acting **strategically.**

This may be a sensible assumption if, for example, all the players in the economy are 'small', and so their actions really do not have that much of an impact on anyone else. Or, in the case of the monopoly if everyone in the economy is small *except for the monopolist,* and so only the monopolist should take into account the effect of their action. However, there are lots (and lots) of cases in which ignoring strategic interactions is to miss out on perhaps the most interesting aspect of the situation. Consider the following examples:

- Dell, HP and Apple are deciding on how to set the prices of their laptops
- Two nations are deciding how much to invest in military equipment
- A penalty taker in soccer is deciding whether to go left or right, and the goalkeeper is deciding whether to dive left or right

- Terrorists are deciding which target to hit, and the Department of Homeland Security is trying to decide which targets to defend
- Your evil professor is trying to decide which questions to set to get you to fail your exam, and you are trying to decide what to revise.

The key thing about all these situations is that they are **strategic,** in the sense that each of the 'players' in the 'game' need to take into account the actions of the other in deciding what to do. More specifically, the best thing for one of the players to do depends on the actions of the other player. Take for example the soccer goalkeeper. If the penalty taker is going to go left, the best thing for the keeper to do is to dive to the left. But then the best thing to do is for the penalty taker to go to the right. But now the best thing is for the goalkeeper to dive to the right, and so on. What would we expect the outcome of this situation to be? None of the tools we have developed so far give us the answer.

The study of strategic interactions of this type, called 'game theory', has been one of the big innovations of theoretical, empirical and experimental economics in the second half of the 20th Century - it is also the only branch of economics, as far as I am aware, whose invention has had a film made about it. It is also, in my opinion, on of the most interesting, powerful and counter-intuitive tools that we have in the economics arsenal.

It is also a lot of fun (how could anything called game theory not be?) In class, we will learn how to answer the following questions:

Example 1 *This story involves a village high up in the Italian Alps. The occupants of this village confirm to all currently available stereotypes. First, the men are Lotharios, in the sense that some of them are cheating on their wives with the wives of other men. Second, they are dreadful gossips, so every man in the village knows whether every* ***other*** *man in the village is being cheated on by his wife (but he does not know about his own wife). Third, they are fiercely proud (and sexist hypocrites) - and each man declares that if he catches his own wife cheating, he will shoot her in the town square at midnight. Fourth, they are very religious, and all attend mass every Sunday. One Sunday, a new young firebrand priest turns up to give a sermon. As part of his sermon he condemns the town as a den of wickedness, with the words "everywhere I look in this village, I see sin. I know for a fact that some of the men in this village are lying with the wives of other men".*

For the first night after the preacher leaves, all is quiet, as is the second night. On the third night, shots are heard in the square at midnight. The question is, how many shots were fired, and how many husbands were cheating on their wives

Example 2 *The beauty contest game: A newspaper runs the following contest: each entrant sends in a number between 0 and 100. Whoever sends in the number that is closest to $\frac{2}{3}$ of the average of all the numbers sent in gets the prize. What number should you send in?*

2 An Introduction to Games: The Prisoners Dilemma

In order to be able to properly analyze situations like this we will need to be a little be a bit more formal about what a game really is. A description of a game consists of three things

1. A list of the players of the game
2. A list of the actions that each player can take
3. For each player, a description of the payoff they get, given the actions taken by themselves and the other players of the game

What makes this situation a game, or strategic, is the fact that each player's payoffs depend not only on their actions but also the actions of all the other players.

In order to illustrate the idea of a game, we will use the most famous example from all of game theory: the Prisoner's Dilemma. This game, has literally spawned thousands of theoretical and experimental papers, so it is a good one to understand.

Example 3 (Prisoner's Dilemma) *Following the great Providence jewelry robbery, Algenon and Betty have been hauled in by the police. Without confessions the police have enough evidence to charge each of them with the minor crime of handling stolen goods, for which they will go to jail for a year. The police offer both Algenon and Betty the following deal. If one of you rats the other one out while the other keeps quiet, then we will let the rat go free, and send the other down for 6 years. However, if you both rat each other out, then you will both go to jail for five years.*

So, how would we formally describe this game? Looking back at the list above, we need to do find three things. First, a list of players. Here, this is simple, there are two players - Algenon and Betty. Next we need a set of actions for each player. Again, in this case, this is pretty simple - each player can either rat (R) or keep quiet (Q). Finally, we need to describe the payoffs of each player, given the actions they both take. Here there are 4 possibilities

1. Both Algenon and Betty keep quiet, in which case they both get 1 year in prison (which we will denote as -1)
2. Algenon keeps quiet and Betty rats, in which case A gets -6 and B gets 0
3. Betty keeps quiet and Algenon rats, in which case B gets -6 and A gets 0
4. Both rat, in which case both get -5

We are going to find it very convenient to represent a game in the form of a matrix that looks like this

		Betty	
		Rat	Quiet
Algenon	Rat	-5,-5	0,-6
	Quiet	-6,0	-1,-1

All the information we need to know about the game is incorporated in this matrix. Algenon's two strategies are represented by the rows of the game (we sometimes call this person the row player), while Betty's strategies are represented by the columns (this player is therefore the column player). The payoffs that each player gets are represented by the numbers in each cell of the matrix - Algenon's payoff first and then Betty's. So, for example, if we want to know what will happen if Algenon rats and Betty keeps quiet, we look along the row in which Algenon rats (the top row) and down the column in which Betty keeps quiet (the right hand column), and see that the payoff for those strategies is that Algenon gets 0 and Betty gets -6.

So, now we have described our prisoner's dilemma game, what do you think the outcome will be. Will either of our protagonists keep quiet? Will either of them rat? In order to answer this question, let's first think about how Algenon might behave in this game. In fact, we should ask the following two questions:

1. If Betty keeps quiet, what is the best thing for Algenon to do?
2. If Betty rats, what is the best thing for Algenon to do?

We call this Algenon's **Best Responses** to each of Betty's strategies.

We can answer these questions very quickly using the matrix we have used to describe the game. First, let us imagine that Betty has decided to rat. This means that we are in the first column of the matrix. Algenon now has the choice to either rat (top row) or keep quiet (bottom row). The top row gives him -5, while the bottom row gives him -6. Thus if Betty rats, Algenon would rather rat. To indicate this, we will highlight Algenon's payoff in the top row in the first column

		Betty	
		Rat	Quiet
Algenon	Rat	**-5**,-5	0,-6
	Quiet	-6,0	-1,-1

What about if Betty keeps quiet? Now we are in the second column, so now Algenon has a choice between 0 if he rats (top row) and -1 if he keeps quiet (bottom row). Thus, if Betty keeps quiet then he is also better off ratting

		Betty	
		Rat	Quiet
Algenon	Rat	**-5**,-5	**0**,-6
	Quiet	-6,0	-1,-1

So, whatever strategy Betty chooses, Algenon would be better off ratting. In this case, we say that to rat is Algenon's **Strictly Dominant Strategy**

Definition 1 *For a particular player, we say that a strategy is strictly dominant if it gives a higher payoff than all other strategies, regardless of what the other players choose to do.*

It seems like a fairly sensible assumption to make that, if a player has a strictly dominant strategy they will choose to take it. Thus, we would expect Algenon to rat to the police.

You should find it easy to check that it is also strictly dominant for Betty to rat. First, assume that Algenon is going to rat, and so we are dealing with the top row. Betty has a choice between ratting and getting -5 (left column) or keeping quiet and getting -6 (right column). Similarly, if Algenon keeps quiet (bottom row), then Betty has a choice of ratting (left column, payoff 0) or keeping quiet (right column, payoff -1)

		Betty	
		Rat	Quiet
Algenon	Rat	-5,-5	0,-6
	Quiet	-6,0	-1,-1

Thus, the outcome we would expect from this game is that both Betty and Algenon will rat, and both will spend 5 years in jail.

This is a result that has absolutely fascinated economists since it was first discovered. Despite the fact Algenon and Betty could spend only a year in jail if they both kept quiet, our prediction is that they will spend 5 years in jail. Put another way, if people act selfishly, the outcome will not be Pareto optimal. This is in stark contrast to the result we got in the case of competitive equilibrium . Moreover it seems to be a result that could be applied much more widely than just to poor Algenon and Betty. For example, here are some other cases that can be fit into the prisoner's dilemma framework:

- You are part of a small tribe which is about to be attacked by another tribe. If everyone fights, there is a small chance that any one person will be killed, but the tribe will fight off the invaders. If no one fights then your tribe will be taken into slavery. However, if you are the only one that fights, then something really nasty is going to happen to you, whereas if everyone else fights and you don't then you will certainly survive, and the invaders will be driven off.

- The USA and USSR are deciding whether or not to nuke each other. If neither country nukes the other, then there is no nuclear war, but they have to continue sharing resources. If both countries nuke, then there will be a war, but neither country will be wiped out. However, if one country nukes and the other does not, then the nuking country gets to take over the world, while the other country gets wiped off the face of the map.

- China and Europe are deciding whether or not to reduce your carbon footprint. If both countries reduce their carbon footprint, then both get a high payoff from the world not sinking. If neither do, then the world sinks, but neither gain a competitive advantage. If only one country reduces emissions, then that country loses all competitive advantage, while the other country gains competitive advantage and benefits from less sinking

What these examples should make clear is that the prisoner's dilemma result raises really fundamental questions about the origins of human cooperation. The truth is that we do see villagers banding together to defeat invading tribes, and we do not see countries nuking each other, and so on. Sociologists, economists, anthropologists and evolutionary biologists have spend a long time trying to square that fact that people DO cooperated with the prisoner's dilemma insight. While a review of this work is beyond the scope of this course, some of the lines of enquiry include

1. Repetition and reputation: Often these games are not played just once, but are played repeatedly with the same people. In such cases, one can make a convincing argument that it is possible for people to coordinated on the cooperative outcome.

2. Punishment. Imagine that both Algenon and Betty are members of the mob. If either of them rat, then they will get their thumbs broken (which they do not like). This changes the payoffs of the game in the following way

		Betty	
		Rat	Quiet
Algenon	Rat	-10,-10	-5,-6
	Quiet	-6,-5	-1,-1

   Now the dominant strategy for both players is to keep quiet

3. Commitment. Imagine that, before being caught, Algenon and Betty wrote a legally binding contract that, if the other person went to jail for longer than they did, they would have to pay the $1 million dollars. This would change the incentives in the same way as the mob contract above

4. Social preferences. A more behavioral explanation is that Betty and Algenon may actually care about each other, and so Algenon prefers outcomes in which Betty does not go to jail and visa versa. In fact, one fascinating line of research is looking at the extent to which the origin of such emotions may be to support cooperation within groups, and break the prisoner's dilemma

3 Iterated Deletion of Strictly Dominated Strategies

Undoubtedly the prisoner's dilemma is a very interesting game to study, but game theory as a subject is much richer than that. In fact, in the prisoner's dilemma, the choices of Algenon and Betty don't really depend on their beliefs about what the other person will do - whatever the other person will do, they are better off ratting. The richness and subtly of game theory really comes from it's ability to handle cases where this is not true - the behavior of Algenon depends on his beliefs about what Betty will do, while the behavior of Betty depends on her beliefs about what Algenon will do.

As an example of this, consider the following modification to the prisoner's dilemma game. Algenon's strategies are the same, but now Betty has a new strategy: She can go Berserk, break out of her cell, and come looking for Algenon. If she does this, then she will be sent down for 50 years, whatever Algenon does. Algenon's payoffs will depend on whether or not he has ratted. If he has, the Betty is going to do something really nasty to him, and his payoff will be -50. If he has, then Betty will break him out of jail, and he will get away scot free.

We can represent this modified game once again using a matrix

		Betty		
		Rat	Quiet	Berserk
Algenon	Rat	-5,-5	0,-6	-50,-50
	Quiet	-6,0	-1,-1	0,-50

What will the outcome of this game be now? We can start by highlighting the best choice for Algenon for each of Betty's strategies

		Betty		
		Rat	Quiet	Berserk
Algenon	Rat	-5,-5	0,-6	-50,-50
	Quiet	-6,0	-1,-1	0,-50

The first thing to note is that Algenon no longer has a strictly dominant strategy. His best action is going to depend on what he thinks Betty is going to do. If he thinks that Betty is not going to go berserk, then he is best off ratting her out. However, if he thinks that she is going to go berserk, then he is better off keeping quiet. Thus, as things stand, we cannot make any predictions about what the outcome of the game is.

What about if we look at the game from Betty's perspective?

		Betty		
		Rat	Quiet	Berserk
Algenon	Rat	-5,-5	0,-6	-50,-50
	Quiet	-6,0	-1,-1	0,-50

Whatever Betty thinks that Algenon will do, she is better off ratting. In particular, she will *never* go berserk. In other words, the strategy 'rat' strictly dominates the action berserk. Thus, it seems that Algenon should put a probability of zero on Betty going berserk. Thus, Algenon should really ignore this possibility when making this decision. In this case, the game looks like it did without this new strategy for Betty

		Betty	
		Rat	Quiet
Algenon	Rat	-5,-5	0,-6
	Quiet	-6,0	-1,-1

Once we have removed this strategy, it is now the case that Rat becomes Algenon's dominant strategy, and this is what we would expect to do.

This process is called the *iterated deletion of strictly dominated strategies.* The idea is that you initially look for dominated strategies for one of the players (Betty), and remove them from consideration - the idea being that the other player (Algenon) should assume that Betty will not play this strategy. We then look at Algenon's strategies in the new game (ignoring Betty's dominated strategies) and see if he has any dominated strategies, we remove them from consideration. We then go back to Betty and see if we can remove any further strategies and so on. Hopefully we end up with a unique strategy for Algenon and a unique strategy for Betty, and that is how we expect the game to play out.

4 Nash Equilibrium

An obvious question is: can we 'solve' every game using the iterated deletion of strictly dominated strategies. The answer is no, as the following game illustrates

		Woman	
		B	S
Man	B	2,1	0,0
	S	0,0	1,2

This game was original called the 'battle of the sexes' game, but has latterly been relabelled as the more politically correct 'Bach or Stravinsky' game. The idea is as follows; A man and a woman are trying to decide what to do in the evening. The man would rather go and watch the boxing (or go and see Bach), while the woman would rather go shopping (or see Stravinsky). However, they would both rather spend the evening together (doing either activity) than they would to do something separate. Thus, their payoffs are as in the above matrix. If they both go to the boxing, the man gets a payoff of 2, while the woman gets a payoff of 1. If they go shopping, then the woman gets a payoff of 2 and the man gets a payoff of 1. If they do different activities then they both get a payoff of zero.

What would we expect to be the outcome of this game? Well, let's first look at the man's best responses.

		Woman	
		B	S
Man	B	**2**,1	0,0
	S	0,0	**1**,2

No dominant (or dominated) strategy here - if the woman chooses S, he would rather choose S. If she chooses B, he would rather choose B. So we cannot delete any strategies here.

What about for the woman?

		Woman	
		B	S
Man	B	2,1	0,0
	S	0,0	1,2

No domination here either! If the man chooses B, she would rather choose B, while if he chooses S, she would rather choose S. So we are stuck! This is a game we cannot solve using iterated deletion of strictly dominated strategies.

So what can we do? One solution was proposed by John Nash - he of 'A Beautiful Mind' fame. And his idea was beautifully simple. He asked the following question: what strategies for each player would represent a *steady state* of the game - i.e. one in which neither player had an incentive to change, given what the other one is doing? Surely we might expect the game to 'end up' at such a point somehow? At the very least, we would expect solutions to the game that do not have this property be unstable? Thus was born the concept of the **Nash Equilibrium**

Definition 2 *A Nash Equilibrium for a game is a strategy for each player such that no player has an incentive to change their strategy, given what everyone else is doing*

So how do we find the Nash equilibrium of a game? One way is to draw the best responses for each player on the same matrix

		Woman	
		B	S
Man	B	2,1	0,0
	S	0,0	1,2

Any cell that is a best response for both players is a Nash Equilibrium. In this case, look first at (B, B). If the woman is playing B, then the man would rather play B than S. Similarly, if the man is playing B then the woman would rather play B than S. Thus (B, B) is a Nash equilibrium. Similar logic will tell you that (S, S) is an equilibrium. On the other hand, (B, S) is not a Nash equilibrium. If the woman is playing S, then the man would rather play S than B. And if the man is playing B, then the woman would rather play B than S.

How should you think about a Nash Equilibrium? Is it a predicted outcome - in other words, is it the outcome that we would expect to happen if the game is played? This, it turns out, a pretty subtle question, and one that has perhaps not received as much thought as it should have done. The simplest interpretation of the Nash Equilibrium is just that it is a steady state of the game - a situation in which all the players of the game hold correct beliefs about what the other players are doing, and act accordingly. Thus the Nash equilibrium of a game is a situation in which no-one has any incentive to change what they are doing.

So do we expect people to end up at a Nash equilibrium? There are good reasons to think the answer is no. Imagine, for example, two strangers playing the battle of the sexes game for the first time, with no chance to communicate with each other. How would they know whether to coordinate on B or S? Even if the man (say) is prepared to give up on the boxing to ensure that he gets to go with the woman, how is the woman supposed to know this? Thus, it is far from clear that we should expect to end up in a Nash equilibrium in this case.

In general, there are two main justifications for thinking that players of games may end up at Nash equilibria.

1. Self-enforcing agreements: If players in the game are allowed to talk beforehand, a Nash equilibrium can represent a self enforcing agreement, in the sense that no one person has an incentive to deviate from the agreement

2. Repeated play: If the game is repeated a number of times, then people are likely to end up playing a Nash Equilibrium.

Unfortunately, there are problems with both these explanations that lie beyond the scope of this course. For now, you should remember two things. First, Nash equilibrium are very important within the study of economics, so it is important that you know how to find them. Second, you should make sure that you take care when interpreting a Nash equilibrium outcome, and in particular you should not confuse it with a cast iron prediction about what is going to happen. You should also be aware that there are lots and lots of game theory papers that look at the circumstances under which people play Nash equilibria.

It may be the case that, even within the same game, there may be some Nash Equilibria that you find more convincing than others. Consider the following game, sometimes called the Mozart or Mahler game.

		Woman	
		Mo	Ma
Man	Mo	2,2	0,0
	Ma	0,0	1,1

This is a variation of the battle of the sexes game. Again, each player would rather spend the evening together, but in this case they *both* prefer Mozart to Mahler. What are the Nash equilibria of this game? Let's look at the best responses

		Woman	
		Mo	Ma
Man	Mo	2,2	0,0
	Ma	0,0	1,1

Obviously it is a Nash equilibrium for them both to go and see Mozart, but it is *also* a Nash equilibrium for them both to go and see Mahler. Most people see this as an unlikely outcome of this game: how on earth would the two of them end up coordinating on Mahler when Mozart is better for both of them (the equilibrium where they both play Mozart is sometimes called payoff dominant - i.e. a Nash equilibrium that pareto dominates all other Nash equilibria of the game).

Note, however, that it is not always the case that a payoff dominant Nash equilibrium is the obvious outcome of a game. Consider the following adjustment to the Mozart Mahler game: Now, (in a surprising decision by the management) anyone turning up at a Mozart concert without a partner is shot, which gives a payoff of -1000.

		Woman	
		Mo	Ma
Man	Mo	2,2	-1000,0
	Ma	0,-1000	1,1

Now, (Mozart, Mozart) is still the payoff dominant Nash equilibrium. However, if either player has even a small amount of uncertainty about what the other player will do, then you are likely to play Mahler. In this case, we sat that (Mahler,Mahler) *risk dominates* (Mozart, Mozart).

4.1 Nash Equilibrium and Firm Behavior

4.1.1 Cournout Competition

We are now going to examine an application of Nash Equilibrium to firm behavior. Specifically, we are going to consider what happens when there are two firms competing in the same market. Note that this is not a situation that we have really studied before. When we talked about perfect competition we did consider the possibility of multiple firms, but they did not really compete with each other, as each firm took prices as given, and so did not think about the way that their behavior affected other firms and visa versa. One justification we used was the idea that each firm was 'small' relative to the size of the market. Of course, when we were thinking about monopolists, the firm was 'large', but by assumption there was only one of it. Here we are going to think about the intermediate case.

We are going to begin by thinking about two firms, who we will imaginatively call firm A and B who sell goods in the same market. Each firm has constant marginal costs c (so the cost of firm A making q_A is cq_A and the cost of firm B making q_B is cq_B). The key to this model is that the price that each firm can charge is a function of *total* output made by the two firms. Specifically, the price in the market is given by $p(q) = a - b(q_A + q_B)$. Thus the profits made by firm A are given by

$$\begin{aligned} \pi_A &= p(q)q_A - c(q_A) \\ &= (a - b(q_A + q_B))\, q_A - cq_A \end{aligned}$$

and for firm B by

$$\pi_B = (a - b(q_A + q_B))\, q_B - cq_B$$

What is profit maximizing behavior for each firm in this case? Let's start by assuming that firm A knows what q_B is, and seeing what their optimal output is (in other words, by calculating A's best response function). We can do this by just treating q_B as a constant, and differentiating the profit function with respect to q_A

$$\begin{aligned} \frac{\partial \pi_A}{\partial q_A} &= a - 2bq_A - bq_B - c = 0 \\ &\Rightarrow \quad q_A = \frac{a - bq_B - c}{2b} \end{aligned}$$

so, crucially, the profit maximizing output of firm A depends on the output of firm B. The more that firm B produces, the less that firm A would like to produce. This is what makes this situation a game. Similarly, firm B's best response function is given by

$$q_B = \frac{a - bq_A - c}{2b}$$

So what would we expect to be the outcome of this game? Firm A could pick a given level of output, to which firm B could respond. But then, in general, firm A would want to respond to B's response, and so on. Where would we expect this to end? This is where the concept of Nash equilibrium can come in handy. The Nash equilibrium is the steady state of such a process - an amount of output for each firm such that the amount that firm A is producing is profit maximizing given what B is doing and visa versa. It seems intuitive that this is where this game could end up.

So how do we calculate the Nash Equilibrium in this game? Well let's say that q_A^* is the Nash equilibrium output for firm A and q_B^* is the Nash equilibrium output for firm B. We know that q_A^* is the best response to q_B^*, and so

$$q_A^* = \frac{a - bq_B^* - c}{2b}$$

Similarly, we know that q_B^* is a best response to q_A^*, and so

$$q_B^* = \frac{a - bq_A^* - c}{2b}$$

This is two equations with two unknowns (q_A^* and q_B^*) and so we can solve for both by (for example,) substituting in for q_B^*

$$\begin{aligned}
q_A^* &= \frac{a - bq_B^* - c}{2b} \\
&= \frac{a - b\frac{a - bq_A^* - c}{2b} - c}{2b} \\
&\Rightarrow \; 2bq_A^* = a - b\frac{a - bq_A^* - c}{2b} - c \\
q_A^* &= \frac{a - c}{3b}
\end{aligned}$$

and similarly

$$q_B^* = \frac{a - c}{3b}$$

When both firms are producing at this level, neither has an incentive to deviate: firm B is optimizing given A's output and visa versa. The profit that firm A is making is given by

$$\begin{aligned}
\pi_A^* &= (a - b(q_A + q_B))\, q_A - c q_A \\
&= \left(a - b(2\frac{a-c}{3b})\right) \frac{a-c}{3b} - c\frac{a-c}{3b} \\
&= \frac{(a-c)}{3}\left(\frac{a-c}{3b}\right) \\
&= \frac{1}{9b}(a-c)^2
\end{aligned}$$

Is this globally the best that the two firms could do? What if they chose to cooperate, maximize joint profit and split that profit between them? In other words, what if the firms operated as a cartel? Joint profits are given by

$$\begin{aligned}
\pi &= p(q)(q_A + q_B) - c(q_A - q_B) \\
&= p(q)q - cq \\
&= (a - bq)\, q - cq
\end{aligned}$$

Thus, the level of q that maximizes joint profits is given by

$$\begin{aligned}
\frac{\partial \pi}{\partial q} &= a - 2bq - c = 0 \\
&\Rightarrow \quad q = \frac{a-c}{2b}
\end{aligned}$$

thus, if the two firms split the output in two, then output is given by

$$\bar{q}_A = \bar{q}_B = \frac{a-c}{4b}$$

and profits are given by

$$\begin{aligned}
\bar{\pi}_A &= p(q)q_A - c(q_A) \\
&= (a - b\frac{a-c}{2b} - c)\frac{a-c}{4b} \\
&= \frac{1}{8b}(a-c)^2
\end{aligned}$$

Thus, if the two firms form a cartel, each firm produces less, as

$$q_A^* = \frac{a-c}{3b} > \frac{a-c}{4b} = \bar{q}_A$$

and makes higher profits, as

$$\pi_A^* = \frac{1}{9b}(a-c)^2 < \frac{1}{8b}(a-c)^2 = \bar{\pi}_A$$

However, this simple example is usually used by economists to exhibit why they do *not* think that cartels are very likely: Because there is no way that either firm can credibly commit to keeping their output at the cartel level $\frac{a-c}{4b}$. To see this, imagine what firm a would do if they knew that firm B was going to produce at this level. We can read the answer off firm $A's$ best response function

$$\begin{aligned} q_A &= \frac{a - b\bar{q}_B - c}{2b} \\ &= \frac{a - b\frac{a-c}{4b} - c}{2b} \\ &= \frac{3}{8}\left(\frac{a-c}{b}\right) \end{aligned}$$

Thus, if firm A believes that firm B is going to produce at the cartel level, then they have an incentive to cheat, and produce more. In other words, the cartel is not self supporting. Put another way, firm A and firm B are trapped in a prisoner's dilemma. Both would do better were they to cooperate, but neither can credibly commit to cooperation. Does this mean we would expect never to see cartels? No! it just means that firms would have to use one of the strategies we discussed above to get around the prisoners dilemma problem (repetition, precommitment, etc)

4.1.2 Bertrand Competition

The model we discussed above, which is called Cornout competition is one in which firms compete in quantities - i.e. each firm chooses a quantity to produce, and the price is determined by the demand curve. A second model of competition, usually called Bertrand competition, is one in which firms compete in prices.

Consider the following (slightly stylized) game. There are two firms in the market, and exactly one consumer who wants to by one item (a widget). Each firm announces a price at which they will sell one widget to the consumer. The consumer will then buy the widget from the firm that announces the lowest price. That firm then produces the widget at cost c and makes a profit equal to the price they sell the widget for, minus the cost c. The other firm gets a profit of zero. If both firms announce the same price, then the consumer has a 50% chance of choosing each of the firms.

What is the Nash equilibrium of this game? The key thing is that, in any Nash equilibrium, the widget will get sold for a price equal to c. we can prove this by showing that in no Nash equilibrium in which the consumer pays more than c. This we can do this by contradiction. Let us assume that we have a Nash equilibrium in which one firm (say it is firm A) is selling the widget for a price $p_A > c$, and so making positive profits. This means that there are two possibilities for the other firm

1. They are setting a price $p_B > p_A$ and so making zero profits for sure
2. They are setting a price $p_B = p_A$, and making expected profits $\frac{1}{2}(p_A - c)$

But in neither case is this firm $B's$ best response to firm A's price. Say, instead, that firm B sets the price $c + \frac{3}{4}(p_a - c)$, or $\frac{3}{4}$ of the way between p_A and c. This is below p_A, and so firm B will get to sell to the customer for sure, and the profits they will make are

$$\begin{aligned} & c + \frac{3}{4}(p_a - c) - c \\ = & \frac{3}{4}(p_a - c) \end{aligned}$$

which is greater than both 0 and $\frac{1}{2}(p_A - c)$. Thus, this cannot be a Nash equilibrium. In fact, as you will see for homework, the only Nash equilibrium of this game is $p_A = p_B = c$

4.2 Mixed Strategies and the Existence of Nash Equilibrium

In all the games that we have looked at so far there has been a Nash Equilibrium, or a set of strategies for each player such that every player is best responding to the strategies of the others. Is it always the case that we can find a Nash equilibrium.

To answer this question, consider the following game

		B	
		Heads	Tails
A	Heads	1,-1	-1,1
	Tails	-1,1	1,-1

This game is called the Matching Pennies game. The idea is that each player, A and B have a coin which they can either put heads up or tails up. If both coins are the same way up, then player

A wins, and player B pays her a dollar. If the coins are different ways up then B wins and A pays her a dollar. Of course, you could use the same game to represent lots of different circumstances. It is sometimes called the Battle game: Players A and B are generals, with B attaching a city and A defending it. B can either attack at point H or at point T, which A can either strengthen the defenses at point H or T. If B attacks where A has strong deferences the A will win, otherwise B will win. More importantly we can think of this game as a model of a penalty in soccer. B is the striker taking the penalty while A is the goalkeeper. B can either hit the penalty left or right, and A can either dive left or right. If A dives the right way then she will save the penalty, otherwise the striker will score.

What do the best responses look like to this game?

		B	
		Heads	Tails
A	Heads	1,-1	-1,1
	Tails	-1,1	1,-1

Unsurprisingly, there is no pair of strategies such that A is best responding to B and B is best responding to A: If A picks heads then B wants to pick tails, but the A wants to pick tails, and then B wants to pick heads, and so on. In fact, the idea of this game is it represents a situation of conflict, where if one player wins then the other loses - this is a class of games called **zero-sum games.**

So this means that this game has no Nash equilibrium? Well, yes and no. It does mean the game has no Nash equilibrium in **pure** strategies. However, consider the following thought experiment: Say that B flips the coin, and so it has a 50% chance of coming down heads and a 50% chance of coming down tails. What is the payoff that A gets from her two strategies? If she plays heads, then there is a 50% chance that B's coin will come down heads and she will win 1 dollar, and a 50% chance that B's coin will come down tails and she will loose \$1. The expected value of playing heads is therefore \$0. Similarly, the expected value of playing tails is \$0. Thus, she is indifferent between playing heads and tails. In fact, the expected payoff of her flipping the coin herself is also \$0, so that is also a best response to B's strategy. But if A is flipping her coin, then B gets a payoff of \$0 from playing heads, playing tails, or flipping the coin. So if both players are flipping the coin, then they are (weakly) best responding to the strategy of the other player - that is there is no strategy that they can pick that will give them a higher payoff. This is therefore a form of Nash Equilibrium. It is what we call a Nash equilibrium in **mixed** strategies.

More generally, a mixed strategy for a player in a game is a probability distribution over their available strategies: in the above example the mixed strategy was to play heads with a probability of 50% and tails with a probability of 50%. The key to the mixed strategy Nash equilibrium is that the *payoff of the strategies that the player is mixing between has to be the same:* Thus, the player is indifferent between playing any of these pure strategies, or mixing between them. If this is not the case, then mixing cannot be a best response - it must be better to play the pure strategy that gives the highest payoff.

If we consider mixed strategies as well as pure strategies, then (almost) every game has a Nash equilibrium - though of course the relevant probabilities may not be 50%.

How do we find a Nash equilibrium in mixed strategies? Let's consider the following variant of the matching pennies game to use as an example

		B	
		Heads	Tails
A	Heads	1,-1	-1,1
	Tails	-1,3	2,-1

Here, it is still the case that A wants the coins to come up the same and B wants them to come up opposite, so there is no equilibrium in pure strategies. However, now the payoffs have changed so that B wins more if she plays heads and A plays tails than she does if she plays tails and A plays heads. Similarly, A wins more if she matches on tails than if she matches on heads.

So how do we find an equilibrium in mixed strategies? Well, the first thing to realize is that if it is going to be a best response for A to mix between heads and tails it must be the case that heads and tails give her the same payoff. If this wasn't the case then her best response would be to play the option that gives her the higher payoff with probability 100%. Thus, let q be the probability with which B plays heads. If A is going to mix it has to be the case that

$$\text{payoff of A playing heads} = \text{payoff of A playing tails}$$

$$\begin{aligned}
q-(1-q) &= -q+2(1-q) \\
2q-1 &= 2-3q \\
5q &= 3 \\
q &= \frac{3}{5}
\end{aligned}$$

So for player A to happily mix between heads and tails, it must be the case that B plays heads with probability $\frac{3}{5}$.This makes sense: given that A makes more money if they both play tails, for her to be indifferent between the two strategies there has to be a higher probability of her winning with heads.

What about player B? For B to be happy mixing, she must also be indifferent between playing heads and tails, so, letting p be the probability with which A plays heads we have

$$\text{payoff of B playing heads} = \text{payoff of B playing tails}$$

$$\begin{aligned}
-p+3(1-p) &= p-(1-p) \\
3-4p &= 2p-1 \\
4 &= 6p \\
p &= \frac{2}{3}
\end{aligned}$$

Thus, the Nash equilibrium in mixed strategies is for player A to play heads with probability $\frac{2}{3}$ and for B to play heads with a probability$\frac{3}{5}$. With these probabilities, both players are indifferent between playing heads and tails, and so mixing between the two is also (weakly) a best response, in that there is not strategy that will give them a higher payoff. It is therefore a Nash Equilibrium for A to play heads with probability $\frac{2}{3}$ and q to play heads with probability $\frac{3}{5}$.

Notice that something a bit strange is going on here: In order to figure out the probabilities with which player A mixes, we look at B's payoffs: in other words A has to play in such a way as to make B indifferent between their two strategies. Does this seem like a sensible prediction of what people will do? This is a question addressed in a classic experimental paper by Goree and Holt called '10 little treasures of game theory and 10 intuitive contradictions'. They observed what happens when people played the following game once

		B	
		L	R
A	U	80,40	40,80
	D	40,80	80,40

where the payoffs are in cents. It should be obvious that this is just a variation of the matching pennies game, so the mixed strategy Nash equilibrium is for each player to play each strategy with a probability of 50%. How do we test this prediction? After all, we can't observe people flipping a coin, or explicitly doing things randomly. However, if all the players are randomizing with 50% probability on each strategy, we should see 50% of the people playing each strategy. And in fact this is pretty much exactly what Holt and Goree observe

			B	
	Exp. Percentage		48%	52%
			L	R
A	48%	U	80,40	40,80
	52%	D	40,80	80,40

So this looks great for game theory. Hurray! However, Goree and Holt realized that this is not a great test. There could be other reasons for people splitting 50/50 between the two strategies (e.g. they had no idea what was going on, and were just acting arbitrarilty). They therefore tested the following variant of the game

		B	
		L	R
A	U	320,40	40,80
	D	40,80	80,40

So now, U looks much better for player A. What is the mixed strategy Nash equilibrium of this game? Well, we know that if people are going to mix, then they have to be indifferent between the two pure strategies. So if player A plays U with prob P, it has to be the case that

$$\begin{aligned} 40p + 80(1-p) &= 80p + 40(1-p) \\ &\Rightarrow \; p = \frac{1}{2} \end{aligned}$$

While if player 2 plays L with probability q, it has to be the case that

$$\begin{aligned} 320q + 40(1-q) &= 40q + 80(1-q) \\ &\Rightarrow q = \frac{1}{16} \end{aligned}$$

Thus, the change in the payoff for player A should only affect the behavior of player B. Unsurprisingly, this is rejected by the data. The results for the second experiment are

			B	
	Exp. Percentage		16%	84%
			L	R
A	96%	U	320,40	40,80
	4%	D	40,80	80,40

In fact, the response of the B players is roughly correct, but the A players overwhelmingly play U, in contrast to the Nash equilibrium. Given this, B should best respond by playing R with a probability 1.

What is going on here? Why does Nash work well in one case and not in the other? One possibility is that in the former case is a zero sum game - what one player wins is what the other player loses. In such games, it turns out that it is quite simple to calculate the mixed strategy Nash equilibrium: it is just the strategy that maximizes the minimum possible payoff that you can receive. So for example, in the first of Goree and Holt games, imagine that player A is playing U with probability p. If player B plays L, then their expected payoff is $80p + 40(1-p) = 40 + 40p$. If player B plays R, then the minimum payoff is $40p + 80(1-p) = 80 - 40p$. If $p < \frac{1}{2}$, then the worst payoff that A can get is if B plays L, and is $40 + 40p < 60$. If $p > \frac{1}{2}$ then the minimum payoff that A can get is if B plays R, and gets $80 - 40p < 60$. However, if $p = \frac{1}{2}$, then the minimum payoff that A can get is 60 (whatever B does). Thus $p = \frac{1}{2}$ is the minimax strategy, and also the Nash equilibrium. Maybe people are just good at playing minimax?

This was a question addressed by Ignacio Palacios-Huerta. First he tested normal folk in the laboratory. In fact, it turns out that normal people are not very good at playing minmax. He then wondered if maybe there were special groups in the population who were good at playing minmax: specifically those who played games for a living. He did in fact find good evidence that both soccer players and tennis players were good at playing minimax (and therefore the Nash equilibrium) in zero sum games - both in their chosen sport, and in the laboratory. This evidence has, however, been called into question in a more recent paper by Kovash and Levitt.

5 Dynamic Games

So far all the games that we have described fall into the category of **simultaneous move** games: Both players take their action at exactly the same time (or at the least, one player has to choose their actions while ignorant of what the other does). We are now going to look at what happens when we allow for games in which people make their moves one after the other: such games are called **dynamic**, or **sequential move** games.

As an example, let's think of a modification to the battle of the sexes game. Now, rather than the man and the woman having to decide simultaneously what to do, the man gets off work before the woman, and can choose to either go to the Bach or the Stravinsky. Once there, he tells the woman where he is, and she can then decide what she wants to do.

The classic way of representing a dynamic game is using a game tree. The game tree for the dynamic battle of the sexes game is shown in figure 1. How do we read this game tree? Starting at the left, we see that the first choice is made by the man - he can either go to Bach or Stravinsky. The woman can then find herself in one of two situations: Either the man is at Bach or he is at Stravinsky. In either case she has to decide whether to go to Bach or Stravinsky. At the end we see the payoffs of each player, man first.

In order to solve this game, we first need to realize that the woman's strategy has now got more complicated. Whereas before she just had to decide whether or not she wanted to go and see Bach or Stravinsky, she now potentially has to decide what to do in the case where the man is at Bach **and** what to do in the case where he is at the Stravinsky. In fact, we are going to demand that the woman tells us what she would do in either of these cases. More generally, a strategy for a player in each game is a list of actions that they would take at *each* node where they have to make a decision. Here, the man only has one node but the woman has two nodes, and so her strategy must include a plan of action at each node.

How does the fact that the game is now dynamic affect what we might expect to happen in this game? A natural expectation might be that the bloke will use the fact that he gets to go first to get what he wants: He goes to see Bach, and then the woman will end up following him, because the alternative is worse for her. This seems like a sensible intuition. Does it match with our concept of equilibrium? In the case of a dynamic game, we can still define a Nash equilibrium as a strategy for each player that is a best response to what the other is doing. And indeed, the case where the man plays Bach and where the woman plays Bach if the man plays Bach and Stravinsky if the man plays Stravinsky (which we will indicate as (Bach, Stravinsky) is a Nash Equilibrium - no player can do better by deviating.

However, consider the strategy pair by which the man plays Stravinsky and the woman plays (Stravinsky, Stravinsky). Is this a Nash equilibrium? Well, if the man plays Bach, the woman is going to play Stravinsky and he will get 0, so Stravinsky is a best response. And the woman is getting the best possible outcome, so she is clearly best responding. So this is also a Nash Equilibrium.

A Nash equilibrium yes, but perhaps one that we don't find very plausible. Why not? Well the woman is getting the man to play Stravinsky by threatening to play Stravinsky if he plays Bach. But do we really believe that is likely? If the man plays Bach, then the woman has a choice between playing Bach, and getting 1, or playing Stravinsky and getting 0. It seems that if she were really faced with that choice, she would go for Bach. In other words, her threat to play Stravinsky if the man plays Bach is **non-credible.**

So it seems that we need to refine our definition of equilibrium in these games, in the sense that Nash equilibrium allows for things that we don't find very plausible (in a way that wasn't a problem in the static games). How can we rule these things out? The answer, as formulated by Reinhard Selten in the 1970's, is the concept of **Subgame Perfect Nash Equilibrium (SPNE)**. This demand that the equilibrium is not just an equilibrium for the game as a whole, but also for each subgame of the game. Very well, you may say, but what the hell is a subgame? A subgame is just the game starting from any node of the original game. So the game that we are discussing has three subgames: The whole game, the game starting when the man has chosen Bach, and the game starting when the man has chosen Stravinsky (Note that these last two games are very boring, as they only have one player, the woman).

So how does SPNE help us solve our problem here? Look at the Woman's strategy: (Stravinsky,

Stravinsky). Is this an equilibrium in both games? The answer is no. In the subgame that starts when the Man plays Stravinsky then the strategy is fine. But in the subgame where the man plays Bach, it is not a best response for the Woman to play Stravinsky. Thus, this strategy is not a subgame perfect nash equilibrium. In contrast, the strategy set in which the man plays Bach and the woman plays (Bach, Stravinsky) is a subgame perfect Nash equilibrium. In the subgame that starts from the man playing Bach, the best thing for the woman to do is to play Bach, while in the subgame that starts from the man playing Stravinsky, the best thing for the woman to do is play Stravinsky. Thus the woman is best responding in every subgame. And given the Woman's strategy, the man's strategy is optimal for the whole game (the only subgame in which he gets to play).

How do we find a SPNE. In fact, it is very easy for finite games - we can use a process called **Backward Induction.** To do this, start at the end of the game (what we call terminal nodes), and ask what is the best thing for the players at these nodes to do. In this case, the game has two terminal nodes: The one where the man has chosen Bach and the one in which he has chosen Stravinsky. In the former case, the optimal thing for the woman to do is to choose Bach, and in the latter it is to choose Stravinsky. Thus, in a SPNE, these have to be her strategies at these nodes. (see figure 3)

If the man performs this mental exercise, he now knows exactly what will occur if he reaches these nodes. Thus he can replace the subgames that start from these nodes with the payoff that will happen as long as the woman plays rationally (figure 3). He can either play Bach, and get 2, or play Stravinsky, and get 1. He will therefore play Bach. Putting this all together gives us the SPNE of the game (figure 4).

Is SPNE a sensible notion? As is always the case, we can think of cases where it breaks down. One famous example is called the centipede game, shown in figure 5. In this game, players play sequentially, and in any round can either choose to carry on the game, or stop the game and take the payoff available at that stage. If they both carry on to the end, they will receive $900 each. But what is the SPNE? Well, in the last stage, player 2 has a choice between 1000 and 900, so will choose 'down' and get 1000. But moving one node back, player 1 is now faced with a choice of 601 and 600. Player 1 will therefore choose down at that node. Repeating this logic, you find that the subgame perfect nash equilibrium involves player 1 going down initially, and receiving $1, when there was $900 on the table. Most people see this as an unlikely outcome, and indeed experiments suggest that this is not what happens when people play the centipede game.

6 Asymmetric Information

The last topic that we are going to cover in any detail is called Asymmetric Information. This is another relatively new idea that revolutionized large swathes of the profession when it was accepted (though this was not immediately. The first paper in this area, called "The Market for Lemons" was initially rejected twice, before latterly earning its author a Nobel prize). The idea is best illustrated by an example:

Example 4 *Imagine that there is a car market with 100 buyers and 100 sellers. There are two types of car: 'peaches', or great cars, and 'lemons' or terrible cars. Buyers are in need of a car more than the sellers, so each type of car is valued more by buyers than sellers, and both groups value a peach more than a lemon. The values that each group put on each car are summarized in the following table*

	Peach	Lemon
Buyer	4000	2000
Seller	3500	1000

Now here comes the twist (ready?): Only the seller knows for sure what type the car is. The buyer knows that half the cars in the market are lemons, but he doesn't know which car is which (to me, this seems like a reasonable situation - at least that the seller is going to know more about what is being sold than the buyer.)

For the moment, let's assume that both the buyers and the sellers are risk neutral. The question we want to ask is as follows: Is it possible for all the cars in the market to be sold? If this were to happen, what is the maximum that a buyer would be prepared to pay for a car? Well, there is a 50% chance of it being a lemon (and so being worth $2000) and a 50% chance of it being a peach (and so being worth $4000). Thus, the risk neutral buyer would be prepared to pay $3000 for the car. But at this price, sellers who own a peach would not be prepared to sell, as the car is worth $3500 to them. Thus, the only equilibrium of this system is one in which only the lemons are sold at a price of $2000. (why $2000? Think of Bertrand competition)

This is a tremendously powerful idea: the fact that there is asymmetric information (i.e. that the seller knows something that the buyer does not) means that good cars cannot be sold - a situation that we call a market unravelling. And note that it is being crucially driven by the asymmetric information. If the seller did not know whether the car was a lemon or not then they would be prepared to sell the car for $\frac{3500+1000}{2} = \$2250$, so there would be some price between \$2250 and \$3000 at which trade could take place.

It also has a number of other applications. One famous one is adverse selection in insurance. Imagine that there are two types of people in the world - high risk and low risk. High risk have a 1/4 chance of getting robbed (and losing \$100), while low risk have a 1/10 chance of being robbed (and losing \$100). For a moment, assume that there are the same number of high risk and low risk individuals, and that both are risk neutal. Now imagine an insurance company that offers to insure people against robbery, but cannot tell the difference between high and low risk types. What is the lowest price that they could sell insurance for and not make a loss? Well, for each person, the expected cost is

$$0.5 \times \frac{1}{4} \times \$100 + 0.5 \times \frac{1}{10} \times \$100 = \$17.50$$

But, at that price, who will buy the insurance? The high risk guy will for sure, but the low risk guy will not. In other words, the low risk guy will get driven out of the market. leaving only the high risk guy with insurance.