[SQUEAKING]

[RUSTLING]

[CLICKING]

JONATHAN GRUBER: So today, despite my admonition, it doesn't look like a very good attendance. So I'm going to give you all the answers to the midterm just to reward you for coming. I'm joking. But I'm glad you decided to come because it's a fun lecture.

And it's very relevant to what we're going to be learning really for the next few lectures, which is thinking about models of oligopoly. Once again, an awesome word, which is really just a description of what most markets in America are.

So we've covered two extremes. We've covered one extreme for markets, which is perfect competition. Doesn't really exist. We've covered another extreme, which is monopoly, which does exist but is still fairly rare.

The more common form of markets is what we call oligopoly, which is when there's a small number of firms competing to provide the good. So it's more than one firm. So it's more than monopoly, but it's less than free entry. It's less than a large number of firms driving profits to zero.

So let's think about the prototypical example of oligopolistic industry, the auto industry. Most cars in America are sold by fewer than 20 companies, but they're sold by, well, more than one. So if you drive on the road, you see cars from primarily 10, but maybe even up to 20 different car companies.

So is that enough to drive us to a perfectly competitive outcome? That is, are those 10 to 20 car companies enough to drive us to zero profits? Probably not. But is that enough that each company can simply set its price and pick its point on the demand curve? Probably not either. So we're somewhere in between. We're an oligopolistic market.

Now, an oligopolistic market, firms can behave two ways. They can behave cooperatively or non-cooperatively. When there's one firm, there's no issue. It's just yourself. When there's thousands of firms, you can't figure out how to get together. But when you have a small number of firms, you can imagine behaving cooperatively or non-cooperatively.

Cooperatively would be to form a cartel. So if you wanted to cooperate, you'd form what's called a cartel, which is a cooperative organization of otherwise competing firms. The classic example of this is OPEC, the Organization of Petroleum Exporting Countries.

This organization formed in the 1970s of the Middle Eastern countries that were responsible for about 2/3 of the world's oil. They got together and said, let's produce cooperatively. And we'll decide how much to produce and what price to set, as if we were a collective monopoly. So in other words, the idea of cartelization is you effectively turn yourself into a shared monopoly with all the benefits we've talked about in terms of profits and things.

So basically, the idea is-- look, what do monopolies do? They get rid of competition. They maximize profits given the elasticity of demand they face. When you form a cartel, you essentially say, look, let's do that. Let's get rid of competition. Let's just create as big a pie as possible. And we'll figure out some way to split it.

Now, of course, the other extreme is a non-cooperative oligopolistic market where firms don't get together and compete-- where firms-- I'm sorry-- don't get together and agree on what to do. And in fact, that's more realistic. And we'll talk about reasons why next lecture. Next lecture, we'll talk about how cartels work and why they don't typically last.

So today, we'll focus on the more realistic case of a non-cooperative. oligopoly, which is when there's a small number of firms in an industry and they can't get together and cooperate, which makes sense. I mean, these are which grew up hating each other, makes sense it'd be hard to cooperate. And we'll talk about next time why that is.

So what do we do in that case? Well, in that case, things get hard. The reason we start with perfect competition monopoly is because they're relatively easy. To think about oligopoly, it gets super complicated. And we have to turn to a new tool, the tool of game theory.

Basically, the idea of game theory is to think of these firms interacting as if they're playing a game. There's n firms in a market where n is small. They're playing a game against each other. And so there's two key pieces of game theory to keep track of. Think about playing a game.

What are two things you want to figure out? The first thing is your strategy. What is the right strategy to win this game? The second point you want to figure out is equilibrium.

When is the game over? When does the game settle into an answer that all parties agree on? With monopoly, there's a set of rules. Once everybody else is bankrupt, you win. That's the equilibrium. So, basically, any game has a strategy and an equilibrium rule.

It turns out defining equilibrium in a game with explicit rules is easy. Defining it in a game where firms are just competing in a potentially bloodthirsty manner is not so easy. And it turns out, there is an enormous field in economics all about how do you think about equilibrium in game theoretic markets? How do you think about when things finally settle down and end up?

And the most famous of these ideas of how to find equilibrium is what we call the Nash equilibrium. Any you guys ever see or read *A Beautiful Mind?* That movie? It's a little bit. Raise your hands. Yeah, OK.

So a bunch of you saw it. So it's about John Nash, fascinating character. He was a mathematician, brilliant mathematician whose fundamental insight-- and if you see the movie, this is so illustrated well in a bar with a group of men are trying to pick up a group of women. And he came up with this fundamental insight that competition can be bad.

Now, if you think about economics, we've learned-- everything we know about competition is good. Competition is good. It delivers better outcomes for consumers. It reduces deadweight loss. But he actually came up with a theory that, actually, in certain situations, competition can be bad and cooperation can be better. And that was his fundamental insight.

And then he went insane, unfortunately, as can sometimes happen. And the movie is all about that. It's a very interesting movie, well acted, and interesting story about this poor guy who had these brilliant insights and then developed mental illness, severe schizophrenia.

So what John Nash defined is the Nash equilibrium. And the Nash equilibrium is defined as follows. Precisely, Nash equilibrium is the equilibrium where no player wants to change their strategy given all the other players' strategies.

So in equilibrium, where no player wants to change their strategy given all the other players strategies. That is a Nash equilibrium. That is formally, if you hold all other players' behavior constant, I can't do any better. When I've reached that point, I'm in equilibrium. We reached a point where I can't do any better holding everybody's strategies constant, then I'm in equilibrium.

And the best example of this is example some of you may have heard about. And you've certainly seen in movies, the example of the prisoner's dilemma. The prisoner's dilemma is a famous movie device, but almost a [INAUDIBLE] example, which is two guys are arrested for a crime. It's not airtight that they did it. They think the two guys did it. To really make it airtight, they need them to admit it.

So let's say in this example-- and we'll give you examples like this. And you'll need to think about them, interpret them. So the example is two guys get arrested. Right now, today, there is enough evidence to send each of them to jail for one year, enough to send them just one year. But if one guy squeals and says, no, the other guy did it, if he turned state's evidence-- we've learned a lot about this in the FBF trial and the upcoming Trump's trial.

The idea is if one person turned state's evidence or squeals, then they get 0. They walk away. And the other guy gets five years. OK, now if they both squeal, they both get five years. so, basically, the situation is if nothing happens, they each get one year. If either tells on the other person, then the person they told on gets five years.

The way we analyze game theoretic situations like this is we write down what we call a payoff matrix. So write down a payoff matrix. So the payoff matrix look as follows. You got prisoner A on the y-axis here. And prisoner B on the x-axis. And you've got four situations.

In one case, let's say prisoner A is silent. He doesn't talk. And prisoner B also stays silent. They don't talk. Then in that case, A gets one year, and B gets one year.

Now let's say that prisoner A stays silent, but prisoner B talks. Prisoner B squeals. Then prisoner A now gets five years, and prisoner B gets zero. Over here, we have what if prisoner A talks? Well, if prisoner A talks and prisoner B remains silent, then prisoner A gets zero and prisoner B gets five. But if they both talk, they both get five.

Questions about how I set up that matrix, given the words, given the description of the problem? You're going to need to be able to take these descriptions out of matrices like this. So any questions about how I did that?

Now what we want to ask is what should each player do? And the cheat in Nash equilibrium games-- the first step is to say, look, the easiest case would be if there is a dominant strategy, if there's a dominant strategy.

Is there one strategy, which no matter what the other guy does you want to do, then that's the Nash equilibrium. That's the outcome that you're going to do. If there's a strategy that's dominant, that's the strategy you're going to play. The best thing to do no matter what the other person does.

Now, if they could agree to cooperate, if the two of them got together in a room beforehand and they trusted each other and they could cooperate, then clearly what is the dominant cooperative strategy? Yeah, they both stay silent. They both get a year.

But let's say they can't do that. They can't get in a room together. They never put them in a room together in the movies. They keep them separate. And they're scummy criminals. They don't trust each other.

So basically, what happens now if your prisoner A and you're not sure what prisoner B is going to do? What should you do if you think prisoner B is going to stay silent? Yeah, talk because your payoff is better. What should you do if you think prisoner B is going to talk? Yeah?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: Talk. Because your payoff-- if you think prisoner B is going to talk, then you should talk. I'm sorry. I wrote I wrote this-- did I do this right? Yep, I did this wrong. That's what I did wrong. It's no better. If they both talk, I say they both get three. That makes more sense.

I set the example up wrong. If they both talk, they both get three. Your answer is still right. But this makes it non-binding. That involves an inequality. This is non-binding. So clearly, prisoner A is better off talking no matter what prisoner B is going to do.

What about prisoner B? What's their dominant strategy. Yeah? Same thing. It's a symmetric problem. These problems won't always be symmetric, but it's a hell of a lot easier when they are.

So you only have to solve for one person. You can assume the other person does the same thing. So prisoner B's dominant strategy is to talk. So we have a situation where both dominant strategies align. Therefore, we have an equilibrium. And that equilibrium is what? It's they both talk. And they both get three years.

Here's what's fascinating. They are worse off than if they both stayed silent. Competition has yielded a worse outcome. We haven't seen that before in this class. Competition, in the sense of trying to do what's best for you at the expense of the other guy, has yielded a worse outcome than cooperation.

That's what's incredibly innovative about-- Nash, even though being a mathematician, won the Nobel Prize in economics in game theory, went from something which was kind of relatively new when I was a grad student in the late '80s to now being a quarter of the introductory training that graduate students, PhD students in economics get. Game theory has become a dominant, so to speak, element of all economics training because it provides such a valuable tool for thinking about this.

So basically, you end up with what's called the prisoner's dilemma, which is prisoners are each better off by ratting on the other. And it ends up they both end up worse off. And this also shows us how to solve a Nash game, a game with a Nash equilibrium, which is you look for each player's dominant strategy.

And you find the point where they align. The point where the two dominant strategies align is, by definition, a Nash equilibrium. Why? Because they're both doing the best they can given the other person.

Now, if the two dominant strategies don't align, there may not be an equilibrium. You can have situations where A's dominant strategy is different than B's dominant strategy. They don't align. Then you might not get an equilibrium. But in this case, they align. In most cases we'll give you, they align. So you get a Nash equilibrium where they both align.

Now this is all cute and stuff. But like this is a class about economics not about crime. But it turns out the same logic applies to play economic games. Let's take the market for soda. And imagine there's two big players, Pepsi and Coke.

And let's imagine there's a $16 billion soda market. Let's further imagine that absent any advertising, they would split it. That with no advertising, Coke and Pepsi would each get 8 billion out of the market.

But they can advertise. Advertising costs 5 billion. So if they advertise, they only take home 3 billion net. However, if they advertise, they drive the other guy out of the market.

So the choice is don't advertise and split the market, advertise, in which case if you advertise and you win, the other guy is gone. And you get the whole market. But if you don't advertise and you don't win and the other guy advertises, you're gone.

So what's that payoff matrix look like? Well, let's write that down. You've got Coke and Pepsi. Coke and Pepsi. Coke cannot advertise or advertise, OK. Pepsi cannot advertise or advertise. If Coke and Pepsi both don't advertise, they each get 8.

If Coke doesn't advertise and Pepsi does, Coke ends up with 0. And Pepsi ends up with 11-- this is wrong in the notes, by the way, Andrew, we should fix this-- which is the 16 that they got from having the market minus the 5 they spend in advertising.

On the other hand, if Coke advertised and Pepsi doesn't, then Coke gets 11 and Pepsi gets 0. And if they both advertise, Coke gets 3 and Pepsi gets 3. So once again, the dominant cooperative strategy is they both don't advertise.

But the dominant non-cooperative strategy in each case is to advertise. Coke is always better off if it advertises than if it doesn't. Pepsi is always better off if it advertises than it doesn't. So they'll both advertise. And you'll end up with both parties being worse off than they just agreed not to advertise.

This is not an irrelevant example. Until about, I don't know-- it's got to be about 10 years ago. Hard liquor, whiskey, vodka, things like that never advertised on television.

Now, as a kid, I thought that was a law, but it turns out it wasn't. It wasn't a law. It was a cooperative agreement among the manufacturer of hard liquor not to advertise on television. They literally had reached the cooperative equilibrium. And then it broke down. And now you see the ads on TV.

It'd be interesting. I don't know what caused the broke down. Actually, Andrew, let's make a note. I should do a case study on. Should I add that? That'd be good to add. Be interesting to think about what caused that. But they used to cooperate, and they don't.

So basically, this is an example of how the prisoner's dilemma can lead to a worse outcome. Questions about that? This doesn't just apply to Coke and Pepsi or prisoners.

Think about Andrew here. Let's say Andrew is in a new relationship, still trying to feel out his partner. And they have a fight. And Andrew's thinking two things can happen. I can go apologize, and I can go eat crow and say it's my fault. Let's say they're both at fault. Andrew can say, I can go eat crow and say it's my fault.

But then if my partner say it's their fault, I'm going to feel like shit. The partner is thinking the same thing. They'd both be better off if they just apologized because they're both at fault.

But they both don't apologize because they don't want to look like the loser, and they break up. Prisoner's dilemma in action in personal relationships. Basically, it's hard to form cooperative relationships. And non-cooperative equilibria can lead to bad outcomes.

Now this is the very, very tip of the incredibly exciting field of game theory. Game theory is a wonderful and exciting field. We have a great course called 1412, which is game theory, which takes this little bit as its first lecture and from there develops a whole semester of game theory. You might say, well, how can you do a whole semester?

Well, let me give you one example. Let me give you one example of how you can solve the prisoner's dilemma without explicit cooperation. How do you solve for some of-- with explicit cooperation? You can do it through playing a repeated game, a repeated game.

So imagine the following. Imagine Coke and Pepsi play this game every year. These is annual numbers. Play the same every year forever. They play this game every year forever.

Now imagine Coke says the following to Pepsi. It's not a cooperation. It's a threat. Coke says to Pepsi, I will not advertise as long as you don't advertise. But the minute you advertise, I will advertise forever. So Coke says to Pepsi-- they start with no advertising. Coke says, I won't advertise if you don't advertise. But the minute you advertise, I'll advertise forever.

Now think about Pepsi's decision. If Pepsi advertises in period 1, in period 1, they get the whole market because Coke is committed to not advertise. They get 11. They get 11. And then Coke says, screw you advertising forever. They get 3 plus 3 plus 3 plus dot, dot, dot, dot, dot.

What if Pepsi doesn't advertise? Then they get 8 forever. You don't need to be a math genius to see this is a bigger number. So it changes Pepsi's dominant strategy. Coke simply issuing that threat changes the dominant strategy. No cooperation here, simply a threat. And yet it enforces the cooperative equilibrium.

But that doesn't work unless the game is infinite. Imagine that we know 100 years from now, the US will ban soda. And Coke and Pepsi know that. They know at some point this game will end. So this market is going to go away. What is Pepsi's optimal strategy in the 99th year? Yeah?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: Because?

AUDIENCE: It's the last year.

JONATHAN GRUBER: It's the last year. Coke can't punish them. It's over. But Coke knows that. So what's Coke's optimal strategy in the 98th year? Advertise. Because they know Pepsi is going to advertise in the 99th year. They might as well get ahead of the game and advertise the 98th year.

You can work backwards to see the game collapses. From year one, no one will play because they know working backwards that as soon as someone commits not to advertise, they're going to advertise, then just going to advertise forever.

So here's a solution that works, but only under particular conditions, which is an infinite repeated game. That's pretty cool. And that is, once again, the kind of things you can think about with the tools of game theory. What 1412 does is show you how to mathematically to prove these things. Think about all the cool games. It depends on who has what information, who has what bargaining strength. There's all sorts of cool stuff that comes into play here. Yeah?

AUDIENCE: [INAUDIBLE] considering you don't know-- it's obviously not infinite. But you don't know that they're going to stop selling soda in 100 years. There's no set limit.

JONATHAN GRUBER: I think that's why I should have saved my whiskey example for this, which is somehow the whiskey companies had decided they were playing this repeated equilibrium, this repeated game, and somehow it broke down. Now, I don't know what broke. Did they suddenly think, gee, people losing their taste for hard liquor, we better get it while it's good? I don't know.

That's a good question. But in reality, sometimes it works, sometimes it doesn't. And that's what is exciting about game theory is to think about deriving the conditions under which it's likely to work and not likely to work. Good question. Other questions?

So that's game theory in a nutshell. What we're going to do now is move from this vague concept of game theory to a particular model, a particular game theoretic model. And that's the model we'll focus on in this class as the model of non-cooperative oligopoly.

But I'm going to tell you, it's one model of an enormous class of models. It's just a model that's particularly convenient to work with. And it gives the intuition without the math being too hard.

So we're going to move to one particular model. But recognize the model we're using is one of a class of models one could use to model non-cooperative oligopolistic equilibria. That may be the most complicated three words I'll use together in this class, non-cooperative oligopolistic equilibria.

So we're going to talk about the Cournot model of non-cooperative equilibria. This is basically a version of the Nash model. So what we're going to do is we're going to take the general Nash model, and we're going to enrich it, saying beyond having two choices, you have a whole schedule of choices. Your choice isn't just A versus B.

But there's a large set of choices you can make, which is true in economic context, usually in A versus B. There's how much you're going to advertise, et cetera, this whole set of choices you can make. But we're still going to consider two firms playing because it gets really messy beyond two firms. Although ones that get 14.12, we'll teach you how to handle that.

So let's imagine the world of flights from New York to Chicago. And let's imagine due to the hub and spoke problem we discussed last time, there's only two firms that fly New York to Chicago, United and American. Obviously, there's more than that, but just to make this example easy. Imagine only two firms that fly to-- these two firms do dominate the market to Chicago. 90% of flights to Chicago are one or the other of those.

So let's say that these two firms are the only firms in the market. And now we want to ask, how does each firm-- knowing the other firm is out there, How does each firm decide on what price to set and how many flights to run? To do that, we are going to set up the Cournot model.

And the Cournot model is going to say we're going to look for the Nash equilibrium. Instead of phrasing in a vague term like dominant strategy, we'll phrase the Nash equilibrium as follows. This is a long sentence, so bear with me.

The Cournot equilibrium is the set of quantities for each firm such that holding other firm's quantities constant, no firm can obtain a higher profit by deviating. Say it again. The set of quantities for each firm such that holding every other firm's quantity constant, no firm can raise their profits by deviating.

So it's a schedule, a set of quantities, a schedule of quantities that maps out your dominant strategy. It's not like there's one dominant strategy. It's a dominant schedule of strategies.

To see that, let's go to figure 13-1. This is an example of the kind figure you're going to become familiar with over the next few lectures. Figure 13-1 shows the market for American Airlines as if it was a monopoly.

Imagine American Airlines had a monopoly. And imagine we're going to write down arbitrarily that the demand curve for flights to Chicago is 339 minus Q, where Q is the total number of flights. That's the demand curve for flights to Chicago, just arbitrarily written down.

And let's say further that arbitrarily that marginal cost equals $147. So it's 339 minus Q is demand marginal cost 147. So what would a monopolist do? Well, good practice for you all for tomorrow night. What would a monopolist do? Monopolist would set marginal revenue equals marginal cost.

Well, what's revenue? Revenue is 339Q minus Q squared. What's marginal revenue? 339 minus 2Q. So a monopolist sets marginal revenue, 339 minus 2Q, equal to marginal cost, which is 147. And you get that the optimal monopoly quantity QM is 96 flights. And the optimal price is $243.

We get the optimal quantity by solving this equation. You get the optimal price by plugging that back into the demand equation. So if this is a monopoly, this is what they're going to do. Note, of course, that price is well above marginal cost because it's the monopoly. There's a markup. So that's what they would do.

And we can see that in figure 13-1 by drawing the monopolist problem. You see the marginal cost is at 147. You see the demand curve is the blue line. The marginal revenue curve is the red downward sloping line, the maroon line, the darker red line.

We find the point where marginal revenue intersects marginal cost, which is at 96 flights, 96,000, whatever. And then to find the price, we read it off the demand curve, which is the 96 flights, the price is 243. And that's the monopoly outcome. OK, questions about that?

Now imagine-- and this is a weird case. But just imagine for a minute somehow through industrial espionage American knew that United-- imagine that United is in the market. It's no longer a monopoly. And imagine American knows United is going to fly 64 flights, 64,000. This is thousands, but it doesn't matter. It's dropped the unit.

I don't know. Somehow it knows they're going to do 64. Well, what should American do, given the knowledge that United's going to do 64 flights? Well, what it should do is it should compute its residual demand.

What's its residual demand? It's the demand that's left over after United does the 64 flights. So what's its residual demand? Well, its residual demand is 339 minus what it's going to fly-- now it's little q because there's two firms in the market-- minus the 64 flights that American is going to provide.

So it's residual demand is 275 minus q sub A. I'm sorry. Yeah, 275 minus Q sub A. That's its residual demand. Well, armed with that residual demand, we can go back and redo the monopoly example. So look at figure 13-2.

American is originally at the outer blue line, D. The fact that United is flying 64 flights shifts it to the inner demand curve, the residual demand curve D prime. But from there, the magic is the same. We compute marginal revenue, which is MR prime. We find where marginal revenue equals marginal cost.

So what's marginal revenue? Marginal revenue is going to be 275 minus 2qA. That's marginal revenue. I just multiply that by q and then took the derivative. Marginal cost is 147.

So the new optimal q sub A is instead of 96, they're going to fly 64 flights. So the new optimal q sub A. Remember now q sub A is the old big Q because there's now two firms not one firm. The new optimal q sub A is 64 flights. And the optimal price is 211.

The price has come down. Why? Because now part of the market's been taken. So if they want to sell into this market, they've got to have a lower price.

So you get the new equilibrium, which is where the maroon line, the MR prime, hits the red line, the MC. That's at 64 flights and a price you read off the demand curve of 211. Questions about that? Yeah?

AUDIENCE: Does it matter what the price of the airline is setting?

JONATHAN GRUBER: No. The question is, does it matter what price they're setting? Actually, it doesn't because your curve here should depend on how many flights they're going to sell. Now, in reality, it would.

But in practice, it doesn't because we're assuming that we're basically-- as long as we know the other firm's conditions, in particular, if we know their marginal costs and things, we're good. Here, we are going to assume that the marginal cost of American equals the marginal cost of United. And that makes the problem easier.

In reality, it'd be more complicated. But as long as we assume the marginal cost is the same, I don't need to know the price. I already know what they're going to set the price at because they're like me. It makes the problem symmetric. The marginal costs are different. It's not symmetric. It gets more complicated. OK, good question.

So basically, in general, of course, American doesn't know how many flights United is going to fly. So what American does is it asks for every quantity that United might fly. How many do I want to fly? It creates what we call its best response curve. And that you can see as the blue line in figure 13-3.

Figure 13-3 is the best response curve. We know two points on this curve. We know if you-- let me back up. What is this graph? This graph has American-- this isn't price. This isn't a demand graph anymore. This has the American quantity on the x-axis and the United quantity on the y-axis. So we're no longer doing a demand curve. This is now a best response curve.

We know already-- I've already told you two points on this blue line. I've told you that if United flies 0, American flies 96. That's the x-intercept of the blue line. I've also told you that if United flies 64, American flies 64. That's the 64, 64 point.

All this line is technically is going through this exercise over and over again for every possible quantity United could fly. So this exercise is saying for every possible quantity United could fly, what is my best response? In other words, if United is going to fly 192 passengers in a quarter, I should fly none.

Now this problem is symmetric. So United's best response curve is just the flip of American's best response curve. They have the same cost structure. Marginal cost is the same. And they face the same market conditions. They're in the same demand market.

So it's a symmetric problem. So United's best response curve is the flip. If American flies nothing, United should fly 96, which is the y-intercept of the red line. If American flies 64, United should fly 64. And if America flies 192, United should fly nothing.

These are the graphical illustration of each firm's dominant strategy. Their Cournot strategy, remember, is asking, what's the best thing I want to do given what the other person is doing? Well, I told you what the person is doing.

So the dominant strategy for American is to plug in every possible thing United could do and then ask, what's the best thing for me to do? United does the same thing. They plug in every possible thing American could do and ask, what's the best thing for me to do? If those two curves intersect, we have equilibrium.

The Cournot equilibrium here is at the point 64, 64. Why? Because that is the one point where given what United's doing, American doesn't want to change its strategy. And given what American's doing, United doesn't want to change its strategy. Yeah?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: What's that?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: Once again, that's why you take 1412. There's all sorts of fun things. What if there's two intersections/ What if there's no intersection? All sorts of fun examples. Take 1412. We'll go through all of them. We only have a couple lectures for game theory here.

So I can't go through all those. But it gets complicated. Let me just say that. It's called a multiple equilibria. Then where you start determines where you finish. In this model, Doesn't matter where you start. There's one equilibrium that's robust. But in a multiple equilibria where you start could determine where you finish.

So in my example with Andrew, in a single equilibrium model, it doesn't matter who started the fight. But in a multiple equilibrium model, you might have very different outcomes based on who started the fight

So basically, that is the intersection. And you can't be-- here's the key point. American can't say, look, gee, but my profit-- I liked it much better when I was doing this. Like I had 96 flights at a higher price.

So I like that much better. Why can't I just do that? Well, they can't Because American won't let them or American can't because United won't let them. They have to account for the other firms response. And that's what's cool about oligopoly.

That's why the game theory because you have to account for the other firms response. That's the graphic. Let's go through the math. Let's do the Cournot math to show you how this works mathematically. Yeah?

AUDIENCE: [INAUDIBLE] firms will always start at equilibrium and stay at equilibrium?

JONATHAN GRUBER: Well, that's a great question. You're asking about market dynamics. We're calling this a steady state equilibrium. Think of this as short run versus long run. In the short run, probably not.

In the short run, you'll buy a certain number of planes. They'll buy certain planes. You'll fuck around for a while. And then you'll say, ah, this is the point at which, in the long run, things are stable. So let's think of this more as a long run versus a short run. Other questions? Good question. Other questions? OK.

So basically, let's look at the math of the Cournot equilibrium. So the first thing we do-- we don't do this, whatever the hell that is. The first thing we do is we solve for each firm's residual demand curve. It's a symmetric problem. I can solve this for one firm. And I know what I'll get for the other firm. I don't have to solve for one firm.

So what is American's residual demand curve? Well, it's 339 minus qA minus qU. I plugged in 64 before because I had some industrial espionage that has told me 64. But in general, I don't know what United's going to do. So I just put in this amount qU.

Now, what is revenue then? Revenue is 339qA-- this is American's calculation of its revenue-- minus qA squared minus qU qA. What is marginal revenue? Well, I differentiate with respect to qA.

This is American's decision. So what is its marginal revenue? It's 339 minus 2qA minus qU. So now I've got a new term in my marginal revenue equation I didn't have before in other examples, the other firm's quantity inches into my marginal revenue.

So we set this equal to marginal cost, which is $147. Once again, symmetrically, I said marginal cost 147 for everybody. Set that equal to marginal cost.

By the way-- let me go back one second. Go back to the figure. Go back to figure 13-1 and 13-2. One thing that's implicit I didn't say is the example average cost also equals marginal cost for the relevant range because marginal cost is flat.

So the relevant range we're examining here, average cost equals marginal cost. Marginal cost can be the same. And average cost could be structured to be different. And then it'll get complicated again.

So really, we're assuming is basically cost structure is the same across the companies, at least the relevant range. So you set that equal to marginal cost. You solve. You solve that, and you get the qA equals 96 minus 1/2 qU.

So now we don't get a deterministic answer. We get that what I should fly is 96 minus 1/2 qU. Guess what. That's the blue line in figure 13-3. Check it out. It's 96 minus 1/2 qU.

Well, it turns out the problem is symmetric. So United has the same best response function. I can do the same math for United. And I get that qU equals 96 minus 1/2 qA. qU equals 96 minus 1/2 qA because it's symmetric.

So guess what. I have two equations and two unknowns. I have two best response functions intersecting or maybe not, two equations, two unknowns. So I can solve this, these two equations and two unknowns. And I can get that qA equals 64. And by symmetry, qU equals 64. So that is how I just mathematically solved for that intersection.

Now, one thing, you might solve this and find no intersection. Two equations, two unknowns. You might get no rational answer. And that would mean the market would just collapse. Or you could get multiple answers, and that's more complicated. But generally, you get one answer in this class at least. And that would be equilibrium. Yeah?

AUDIENCE: [INAUDIBLE] mathematically, but like conceptually, [INAUDIBLE] qU [INAUDIBLE], assuming there's one monopoly?

JONATHAN GRUBER: Great question. Does anyone know? Great question. Because what's big Q here? Big Q now in the market is 128.

And what's the price? Well, we know the demand curve. The demand curve hasn't changed. The demand curve is 339 minus Q. It's one market demand curve.

So the new price is 339 minus 128, so 211. But Q is 128. [INAUDIBLE] is 96. Great question. Why? Yeah?

AUDIENCE: Because [? I ?] believe the firm that controls everything is a price setter versus an oligopoly. They're price takers.

JONATHAN GRUBER: Well, they're in between price takers and price makers. But effectively, the way we model it, they're price takers. But the intuition is when I'm a monopolist, I am maximizing the profits for me, finding out what maximizes my profits, which is the high price and a low quantity.

Once I'm sharing the market, think about the monopoly poisoning effect. This is a great hint for what we'll cover next time. Think about what happens to monopoly poisoning effect. If I sell more in the market, who bears that poisoning effect? Both of us. Because once Q goes up, once big Q goes up, the price falls. We both suffer in the lower fall.

So in some sense, I'm offloading half of the poisoning effect. When I'm a monopolist, every additional unit I price I want to sell lowers my profit. When I'm in this larger group, there's a sharing of the poisoning effect. So you end up selling more.

What we'll show next time is why exactly this intuition-- it'll be clearer next time-- leads to cartels not to be functional. And it's why we'll show next time the limit of oligopoly is perfect competition. But the bottom line is when you have a Cournot equilibrium, you end up with an answer that is not perfect competition.

Remember, what would perfect competition say the quantity would be? Can anyone tell me? What would the quantity be under perfect competition? What would big Q be in this market with perfect competition? It involves one equation, solving one equation. Yeah?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: And how'd you get that?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: No, if perfect competition. That's exactly this question here. If there's perfect competition, what would-- yeah, in, the back.

AUDIENCE: 64.

JONATHAN GRUBER: Well, it's 64. How'd you get that?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: Well, OK, that's one way to do it. What's another way to do it? Well, if it's perfect competition, what do we know? Come on, guys, got to be ready for tomorrow night. What's the profit maximizing condition of perfect competition? Price equals what?

AUDIENCE: Marginal cost.

JONATHAN GRUBER: Marginal cost. So we just say, well, look, you want to [INAUDIBLE] competition, price equals marginal cost. So what's the price? What's marginal cost? 147. So price equals 147. And then you simply know that the quantity is 339 minus 147.

The quantity is going to be 339 minus 147, which is going to be 192, right? Is that right? No. Oh no, I got this wrong. 339 minus 147-- that's a hard math. 340, so 200. 200.

Yeah, 192, right? Yeah, 192. I'm really insecure about my math skills. So 192. 192 is much bigger than monopoly, which was 96. It's also bigger than what we got in the oligopoly case.

So what we see is oligopoly leads us to be in between monopoly and perfect competition. And we'll come back next time to review that and show you why that is. OK, let's stop there. Good luck tomorrow night. I'll see you all bright and early on Wednesday, and we'll continue this discussion.