[SQUEAKING] [RUSTLING] [CLICKING]

IAN BALL: So I know last class was quite theoretical. It was kind of on the mathier side. So I hope people who like math enjoyed that. Today, the lecture is going to be a bit more on the economic side, so a little less math. I mean, there'll be some algebra, but it's a little more applied, a little more practical.

So today we're going to talk about cartels. And I think it's helpful to start with some examples. So the classic example of a cartel would be OPEC. OPEC is the Organization of Petroleum Exporting Countries. And what OPEC does is they get together, and they try to coordinate the production of oil to ensure that oil prices are high, and therefore, that all these countries can make more money.

So the general idea of the cartels we're going to look at today, maybe I'll say the goal-- there's many different kinds of cartels. So I'm going to talk about the goal of these specific kinds of cartel-- is generally to coordinate on lower production in order to raise market prices.

Do people know any other examples of cartels that maybe you've come across in your experience or in other econ classes, or just things you've heard of? Any ideas? Yeah?

STUDENT: Are drug cartels--

IAN BALL: Drug cartels. Great example. So drug cartels often maybe interact a little bit differently. But they're certainly-- drug cartels are interested in controlling what's on the market and controlling how much revenue they get. Often, drug cartels maybe operate-- maybe the way that they limit production or limit competition is often they segment the market.

So it's common for one drug cartel to, say, operate in the north of the country, another cartel to operate in the south of the country. And they have this agreement that only one cartel is going to produce in one side of the country, and only the other cartel will produce in the other side. And that's one way of reducing production because you only have one side selling in each market. Yes?

STUDENT: What if a new group comes in?

IAN BALL: Right. That can often lead to violence. And even without new groups coming in, there's often a fight over turf. So I like to use this example. This is a good example of a cartel.

It's not really going to fit in with our models because our models aren't going to have the threat of violence. So I think this is an important economic question, how criminal organizations use violence to sustain cooperation. But we're going to think of cartels that threaten to produce more of the good, not to kill the other cartels. So this has violence, which will not really be covered in our models, but it's a great example of a cartel. Other examples of cartels that people have come across?

So there's a famous one, the famous vitamin cartel. In the US, these producers of-- it's maybe not the sexiest industry, but these firms that produced various vitamins, they colluded to try to keep prices low. I think they also split up exactly which vitamins they each produced.

One that I think maybe people don't think of as much that I think is a nice example is actually the production of doctors. So I don't know if any of you are applying to med school. The number of slots in the entering classes at med schools is very restricted. Medical schools cannot enroll more students than they've agreed to with various organizations, like the AMA.

Why do you think doctors have reached this agreement that med schools cannot enroll too many students in the United States? Why might they want to do that? Yeah?

STUDENT: To be able to [INAUDIBLE] salaries.

IAN BALL: Exactly. So it's exactly the same idea. OPEC wants to underproduce oil so that the price of oil is high. How do we produce doctors? Doctors are also something that we produce and sell. I mean, it's a bit weird to think about selling your labor, but that's what's happening.

Well, the production of doctors is med school. And if medical schools reduce the number of doctors they produce, then the price you can sell a doctor for, that is, the wage that doctors get, will be higher. And that's exactly what we see here. And classical guilds also fall into this. Maybe some people would say that the structure of PhD programs also fits into an academic cartel.

And then maybe one last example I'll give that also doesn't quite fit with our model today is actually the Quebec maple syrup cartel. Probably not well known to people. This one is different because this is a government-sponsored cartel.

So the government of Quebec and the government of Canada explicitly allow this cartel because they think that producing maple syrup is a valuable cultural practice in Canada, and they want to sustain the profits and the revenues of producers of maple syrup. But because this is explicit, it's contractual, it has explicit government support, this also won't really fit within our models today.

So we're going to think about models that have no violence, and also no explicit government endorsement. So these are implicit cartels. And I'll say more about what that means.

So the goal of these cartels tells is to coordinate on lower production in order to raise the market prices. What is the challenge? Suppose you're an oil-producing nation. You get together with the other oil producers, and you propose a plan.

We're each going to produce less in order to keep the market price high. That sounds great. Everyone does better. The prices are high. That's good for everyone. What's the challenge here? Why might this not be so easy? Yeah?

STUDENT: Defectors.

IAN BALL: Defectors. And how might they defect?

STUDENT: They start selling more than they agreed to because the price is high, so they make more profit.

IAN BALL: Exactly right. So collectively, this is good for us. But the problem is that each firm or each country will have an individual incentive to deviate by overproducing. So the challenge is that there's an individual incentive to overproduce.

And we can frame this actually in the classical economic language of externalities. If a single firm unilaterally deviates and overproduces, they impose an externality on the other firms. What is the externality that I impose on the other firms if I produce more? How do I affect the other firms when I produce more? Yeah?

STUDENT: Your increased production leads to a lower price than what they can sell.

IAN BALL: Exactly. So when I raise my production, I only internalize the effect it has on my profits. But I'm actually hurting the other members of the cartel because by overproducing I'm lowering the market price, and other people feel that. So this is a classic issue of people don't internalize their externalities, and they overproduce. You can go back to pollution or congestion on roads, all these classical externality questions, this is what's going on.

So what's the solution from the-- I want to say solution from a socially efficient standpoint, but a solution from the cartel standpoint. So I don't use solution with necessarily positive connotation. Well, there's a few solutions. One is to threaten violence, but we're not going to talk about that today. One is to write contracts.

What we can say is, we're going to write a contract that this is exactly how much oil you have to produce, or this is exactly how many units of vitamins you're allowed to produce. But often, we don't see contracts written like this. Does anyone know why? Why do we not see contracts as the solution to this kind of problem? I think there's a few reasons. Yeah?

STUDENT: I think it's probably frowned upon, as [? making ?] [? a ?] cartel. Usually, the government would want to force some kind of fair competition.

IAN BALL: Right. So it's not only frowned upon, it's illegal. It's very frowned upon. So in a domestic context, I think there's two reasons we don't see this. So maybe I'll say but-- maybe I'll start a new board here. But what's the issue?

Well, within a country, this is often illegal. I mean, it depends on the country. But certainly in the United States, if two firms get together, and they were to write a contract that says, we're only going to produce this much, that would go against antitrust regulations, and they'd get sued.

What about-- but I think there's slightly a different reason. What about across countries? So what about OPEC? It's not necessarily legal if Iran and Saudi Arabia sign a contract. It's not really clear what court would say that they were violating antitrust law. But I think there's another problem. So why is it hard for countries to write contracts with each other? Let's see. Yeah? [INAUDIBLE].

STUDENT: Is it like an enforcement issue?

IAN BALL: Yeah. So say more about that.

STUDENT: Like, you have a contract. Like, the only way that you can enforce it is through your own oil production. But like, there's knowing you can snitch on and say that they overproduced.

IAN BALL: So I think you're getting at two things. One is I think what you're getting at is kind of observability, that maybe it's hard to see who's overproduced. I guess I was putting at-- but your word "enforcement"-- I think one basic question is, who enforces contracts between countries? If there's two firms in the US, and they write a contract, and someone breaches the contract, they can go to court in the US.

If Iran and Saudi Arabia sign a contract, well, who do they go to? It's not clear who can enforce that contract. So across countries, it's not clear how we enforce contracts, kind of no enforcement mechanism-- no enforcement, maybe I'll say, through courts.

So in both the domestic context, where collusion as a cartel is illegal, and the international context, where enforcing contracts between countries is challenging because we don't really have a court to enforce them, our initial solution of writing contracts becomes a challenge. And because of that, what we often turn to are implicit cartels.

So by an implicit cartel, I mean we're not writing down a contract that says, if you produce this much, this is what happens. But we're going to try to arrange our behavior in such a way that if someone were to overproduce, they would be punished, maybe by the other firms producing more as well to punish them.

The problem is we have to make sure that this is self-enforcing. But it must be-- maybe I'll say self-enforcing. So the idea is we don't have an external court to take you to jail if you violate it.

But now, well, if someone overproduces, then they're going to be punished by another firm. But then we have to make sure, is that other firm going to want to actually punish them? Are they going to want to overproduce? And then we might need someone to punish them and so on.

So what we want is exactly a Nash equilibrium. In the absence of courts and contracts, we want this agreement to be self-enforcing in the sense that no one can profitably deviate from this. So what have we reduced the problem to? Well, this is precisely a repeated game between strategic players. And we want to try to understand the Nash equilibria or the subgame perfect Nash equilibria of this game, because that will precisely describe what can be achieved through a self-enforcing interaction.

So today we're going to look at a model of this, which is going to be repeated Cournot competition. So formally, the model today is going to be infinitely repeated Cournot-- remember, Cournot refers to quantity-- competition. You can think about a very abstract cartel, but I think this is a pretty good model for OPEC. So keep that in mind as you're running example as we go over this.

And we want to formally examine the issues we've brought up. We want to see-- we want to formalize this intuitive idea that firms will have an incentive to deviate and overproduce. And then we're going to see how we can try to dampen that incentive through future threats to punish by overproduction. So let's get into that.

So what's the model? The model in each period is just going to be standard Cournot quantity competition. So let's say we have firms, i equals 1 through n. And we have periods. Well, 0, 1, 2, 3, and so on.

Again, there's no natural point in time when OPEC breaks down. Maybe eventually we run out of oil, but we don't know exactly when that is. So I think the infinite horizon model is a pretty good model for that. And to simplify things, we're going to assume that each firm has constant marginal cost production.

So C, where we have 0 C less than 1. This is 1.1 stylization of the model. In OPEC, not all firms are probably symmetric. Some countries have better resources to produce oil. They can produce oil more cheaply. You could extend the model to an asymmetric model. But to keep things simple, we're going to imagine that each of the firms involved has the same marginal cost.

And then demand is going to be, as we always say, the standard demand model, we'll write P of Q equals the max of 1 minus Q and 0. Remember here, P of Q is the market price that's faced by all the firms. So this would be the market price for oil.

The market price for oil depends on the quantity of oil that's on the market, and it's decreasing in the quantity of oil on the market. The more oil is on the market and is produced, the lower the price is. The price can never go negative. So that's why we have the max with zero.

So if tons of oil is produced-- well, there was a time recently actually where the price of oil got negative, but it was only a little bit negative because no one wanted to store it during COVID. But in general, the price of oil is generally not going to be negative.

And how do we interpret this one-- I mean, one should be like whatever the total amount of oil that could possibly be produced in the world is. So one is a very big amount of oil. We're just changing our units to make that one unit of oil.

And Q, remember, is going to be q1 plus qn. So we're going to use the notation from before, that little q's refer to how much is produced by a single firm. The big Q is the total amount produced by all the firms together. So the total amount produced is the sum of the amounts produced by each of the n firms.

And then we're going to use standard discounting. So everyone's going to have the same discount factor delta. So what that means-- maybe let's try to write this out formally.

So what I want to write down now using this discount factor is, let's suppose that this is the sequence of production levels. Now, there's no subscript here. So what Q superscript 0 means is the vector of quantities produced in period 0.

So remember, superscripts are always about time, and subscripts are always about the firm. If there's no subscript, this should be interpreted as a vector that tells me how much every firm produced in period 0. Then q1 is another vector that tells me how much every firm produced in period one. And I want to compute here what is firm i's average discounted profit if this is the sequence of production levels?

Well, it's going to be a summation. We're going to use 1 minus delta up front, like we always do, so that we look at average discounted profits instead of just the sum of discounted profits. And then we're going to take a sum over all the periods, so from t equals 0 to infinity.

And then we're going to have delta to the t times the profit that firm i makes in period t, so times pi i. Maybe this is little pi i, and this is big pi I. I don't know if you can really tell the difference. And here, we're going to have qt. Remember what is qt? qt is q1t up to qnt. It specifies how much each firm produces in period t.

So we can now say, what is a firm's average discounted profit from a sequence of production decisions? We can express it like this, but we still have to fill in what the per period profit is. So let's recall where, in general, pi i of q1t up to qnt is what?

Well, this says in period t, what does firm i get? Well, we need to, as usual, keep track of what is their profit per unit, and then how many units do they produce? So the number of units they produce is qit.

In the t-th period, firm i is producing qit units. And their profit is going to be that number of units times the profit they're getting per unit. So what is the profit per unit here? Yeah?

STUDENT: [? Will ?] [INAUDIBLE] get the price by summing all the quantities at this time.

IAN BALL: Exactly.

STUDENT: And then use pi C.

IAN BALL: Exactly right. So that's going to be P of q1t-- try to write small here-- up to qnt. This is the sum of the quantities that are produced at time t. So that's the market quantity. That's the amount in the market. That determines the market price, P. But of course, that's not all profit for me because I also have to pay a cost of production. So we subtract C here.

So now we've formalized the game. I guess one final assumption we have to keep track of is remember, whenever we've studied infinitely repeated games, we've assumed that past actions are observed. And in this context, that may be a strong assumption, and this goes to a question we had. Remember, we're going to maintain the assumption. So I'll say maintained assumption past actions are observed, which in this case, past actions are just quantities produced.

Do you think this is a good assumption in this model or not? Yeah. Yeah? I think it's not great. Not terrible. It depends on the context. Sometimes we might be able to see how much people are producing. Sometimes it may be a bit hard.

And if you read when OPEC has disagreements, this is always what happens right. One country says, oh, you definitely overproduced. They said, oh, no. It wasn't me. And you'll see cartels-- I mean, if you watched *Breaking Bad,* this is often a story. People are trying to secretly sell their product without other members of the cartel observing that they're selling it.

And I think maybe a better model of this would be that we observe noisy signals of how much people have produced. Because really, what we observe is the price. And what happens with OPEC is they see the market price has dropped. Everyone sees that, and that suggests that someone has overproduced. The hard thing is to figure out who it was who overproduced.

And that's an exciting topic, kind of getting more modern research ideas, what we do when we only have imperfect monitoring of past quantities. Unfortunately, it's mathematically a lot harder, so we won't cover it in this class.

But if you continue studying game theory, I think that's a really interesting topic. What happens if we see price as gone-- is really low, and we're trying to figure out who it was who did the overproduction, and we don't really know who it is. But we'll maintain that assumption today.

So before we get into actually solving for a subgame perfect Nash equilibrium of this game, let's try to just think about this basic challenge of each firm having an incentive to overproduce. So let's just first think of the static-- let's just first look at the stage game.

And let's see if we can remember-- the stage game has a unique Nash equilibrium that we solve for. I think this is a good thing to keep in mind maybe for exams. The stage game Nash equilibrium maybe I'll say was q1 equals qn.

Remember, now, we're just thinking of the stage game. So let's just think about everyone's choosing how much to produce, and everyone produced maybe what I'll call qNE. This is the Nash equilibrium quantity. So this is a number. And that number turned out to be 1 minus c over n plus 1.

So it turned out that in equilibrium, the amount that everyone produced-- well, it depends on how many firms there are. The more firms there are in the market, the less each firm produces, and the higher the cost c of production, the less each firm produces.

And what that meant was the market price, pNE, was what? Well, it's 1 minus the total quantity. So maybe let's keep track of this. What is qNE? So I've told you what each individual firm produces.

qNE is the aggregate quantity produced. Well, that's just the sum of the amounts that the individual firms produce. So I have 1 minus c over n plus 1 summed over the n firms. What is that? That's just n over n plus 1 times 1 minus c. I've just multiplied this individual production level by the total number of firms, n.

And once we see that, we can see the market price, pNE, is going to be 1 minus qNE, which is 1 minus n over plus n plus 1 times 1 minus c. And there's a nice-- just to help us understand this, this is kind of a convenient algebra trick.

Let's add c and subtract c. So if we do that, I'll add c to the front. And now I'm going to subtract c later on. So we get 1 minus c minus n over n plus 1, 1 minus c. So all I've done is I've put in a c here, and I've subtracted a c here. So I haven't done anything.

But now we can simplify things a lot because I have 1 minus c minus n over n plus 1 times 1 minus c. So if I factor out the 1 minus c, I get c plus 1 minus c over n plus 1. And I think this is a bit more easily interpretable.

Can people-- does anyone want to give an interpretation for this market price here? So how much is each firm charging in this equilibrium? Maybe I'll just go-- they're charging-- well, their cost plus a markup.

So this is the classic markup term. This tells me that each firm is making positive profits because each firm is producing a positive amount, and the price they're getting is strictly higher than the cost of production. So that means their profit per unit is strictly positive. And the scale of the markup is decreasing in the number of firms. The more firms there are, the smaller this markup is.

And now, let's just compute one last thing, which is what is pi NE? So what is the profit that each firm makes in this equilibrium? Well, each firm is producing 1 minus c over n plus 1. And their profit per unit is exactly the markup, which is 1 minus c over n plus 1. So it turns out this is 1 minus c over n plus 1 squared.

So that's the Nash equilibrium. Now, the question is, how could they do better in the stage game? So I'll ask can firms do better than the Nash equilibrium? This may not be a Nash equilibrium, but I'm just saying, is there another way they could produce to do better than the Nash equilibrium?

And the idea is, so let's say suppose the firms wanted to make the sum of all their profits as high as possible. What would they do? So let's be a bit more specific. How they do better? Well, what they could do is maximize the sum of their profits.

How would you do this? How would the firms maximize the sum of their profits? Yes?

STUDENT: You could choose exactly one firm to produce everything. And the reason is, in that case, n is 1, so then the market can be higher.

IAN BALL: That's right. So I think we're getting very close. Yeah. They want to think about-- so if only one firm produced, how much-- what would be the right amount for that firm to produce?

STUDENT: Then that firm would produce 1 minus c over 2.

IAN BALL: Right. Exactly. And we can actually-- there's a few different ways we can achieve that. So the key is, if we want to maximize the sum of our profits, what we want to do-- and this is-- what you described as a special case of this-- we want to collectively act as a monopolist. We want to make sure that the sum of the amounts we produce is exactly the amount that a monopolist would produce because we know that a monopolist maximizes profits.

So to maximize the sum of the profits, what we need is we need q1 plus qn to equal the quantity that maximizes the sum of the profits. It turns out that's exactly the monopoly quantity. So we want this to be the monopoly quantity. Why is this?

Well, the trick is if I produce q1 and firm 2 produces q2 all the way up to firm n, producing qn, the sum of the firm's profits is exactly the same as the profit that a single firm would make if they produced q1 all the way up to qn by themselves. And the way that we maximize that amount is exactly by setting it equal to the monopoly quantity, which, in this example, is 1 minus c over 2.

So what you described is one special case of this, where one firm could produce the monopoly quantity and everyone else could produce zero. But I don't think that would be a very effective cartel. If I said, hey, other countries, I want you to all to produce zero, and I'm going to produce everything, I don't think those other countries would be very happy.

So probably we would split it a bit differently. A natural split would be an even split, where each of us produces 1 minus c over 2n. So one thing we could do-- so, e,g., we could do q1-- I'll go down here. The even split would be we each produce maybe what I'll call qM.

So before, I described the Nash equilibrium quantity. Now I'm going to describe-- when I say monopoly quantity, what I really mean is each firm's share of the monopoly quantity. So this is 1 minus c over 2n. So qM is technically 1 n-th of the monopoly quantity. Because it's the quantity that each firm needs to produce to ensure that the total production is equal to the monopoly quantity.

Now, if we compare these two numbers, what do you see? Which is bigger, and what's your intuition for that? Let's assume n is at least 1-- at least 2. So which one is bigger and why? Yeah?

STUDENT: The Nash equilibrium is larger. And the intuition for that could be that in order to do better than Nash equilibrium, the firms kind of need to have a closer collaboration to maximize as if they were a single firm. Whereas, each of them in that non-Nash equilibrium state, a monopolist state, would be kind of incentivized to produce more to capture profit.

IAN BALL: Exactly right. So this is exactly what we talked about at the beginning. In the Nash equilibrium, the firms are behaving optimally given the way everyone else is behaving. But that actually turns out that collectively it hurts them because they're producing too much. They would actually do better if they each produced less.

If they each produce less, then the market price is going to be higher, and the payoff of each firm is actually going to be strictly higher. The problem is this will not be in equilibrium. And what will happen is suppose we did this. Let's just check. This is not going to be a Nash equilibrium. And why? What is the deviation for, say, firm i?

What would they-- how could they do better? Would they want to increase or decrease their production? Increase their production, right? So it's not a Nash equilibrium because firm i prefers to strictly increase production.

And we can check that mathematically. But let's just-- this seems a little puzzling now, and this is getting at a core game theory idea. We said that these production levels maximize the sum of everyone's profits. But now I'm saying one firm can do better by strictly increasing how much they produce. So how is that consistent? I thought we were already maximizing profit. So how can one firm do better by deviating? Yeah, on the front?

STUDENT: You're maximizing the [? sum ?] [INAUDIBLE].

IAN BALL: Exactly. So what must happen? If this firm produces strictly more, this firm is doing better. But what do we know must also happen? The other ones are doing worse. In fact, the amount by which they're doing worse more than offsets the amount that this firm is doing better.

So I benefit myself. I hurt everyone else. On average, the sum of payoffs goes down because I hurt everyone else more than I help myself. But I don't care because I'm helping myself. That's the basic intuition of this.

So now the question is, can we sustain the monopoly levels in equilibrium? So the question is-- here's a question. So we're all better off if every period we produce qM rather than qNE. That would be great for us. It's not clear, though, how we can sustain this in a self-enforcing way because we already said if we try to do this, one firm can strictly benefit by deviating and producing slightly more. So how can we do this?

Well, what's the idea? We need some way to punish people for overproduction. And, in fact, we'll just punish them if they produce anything other than what they're supposed to. Really, overproduction is the tempting thing to do. But here's the idea. We want to punish firms for overproduction.

So any ideas about how we might be able to do this? So we're getting together. We see that one firm has deviated. Say, they've produced too much. Now, what might we want to do to try to punish them? Yeah?

STUDENT: We could all move to the Nash equilibrium quantities.

IAN BALL: Exactly. So I think the first-- that's exactly where were going. I was expecting a different answer. So well done. You jumped ahead. What I would expect people to say is we would just flood the market.

The problem with that is that may not be self-enforcing, because why does the firm actually want to flood the market? So indeed, what you're suggesting-- I mean, we could think about doing that later. But indeed, the natural thing to do is, the way we could punish is we could revert to the Nash equilibrium.

Remember, last class when we talked about the folk theorem, we described these strategies where if anyone deviated, we reverted to the Nash equilibrium. In this case-- so by reverting to qNE. So what this says is we're going to try to produce the monopoly level.

We're each going to produce a small amount. We're going to get high profits because the price is very high. If we observe anyone deviating, well, now we're all going to produce the Nash equilibrium level. We're all worse off, but we're actually willing to carry out this punishment because this is itself a Nash equilibrium. That's the idea. And now we want to formally check this.

I guess first question, will this always work? What are we going to need for this to work, for this to be a subgame perfect Nash equilibrium? What's going to need to be true about delta? Yeah?

STUDENT: It needs to be sufficiently large enough that we kind of value the profits in future periods.

IAN BALL: Yeah, exactly. So our guess-- I mean, we'll formally check this, but it's always good to make a guess in advance. Our guess is that this will work if delta is high enough. Because the idea is I can benefit today by overproducing.

The only reason I'm not going to do that is I'm afraid of this punishment in the future. But I have to be patient enough that I care about that future. If I only care about profits today, then, of course, it's not going to be a Nash equilibrium because all I care about is the benefit that I get today.

So now we're going to actually construct this equilibrium. So let's start here. All right. So now let's construct the SPNE that has this form. It's going to be helpful to define two functions, just because if we don't define these, the algebra gets really complicated. So let's just make two definitions-- f of q and g of q.

So for any quantity level q, f of q is going to say if we all produce q, what is our path? So mathematically this is q times 1 minus nq minus c. And technically, this only makes sense if q is less than 1 over n, so that this number doesn't-- so prices don't become negative.

So this says if we all produce q, the market price is going to be 1 minus nq because nq is the market quantity. So 1 minus nq is the market price. And if we subtract c from that, we get the profit per unit, and then we multiply it by q.

And then the other thing we want to understand is how well can a firm do if everyone else is playing q, and I choose my best deviation in the stage game? This is going to be the key deviation we have to worry about. So we want to have notation for this.

So this is going to say, what happens if I choose my quantity level q prime to maximize my profit given that all the other firms are choosing q? Well, if I choose q prime, this is how much I produce. What is the price going to be? Well, it's going to be 1 minus n minus 1q minus q prime, and then I subtract c.

So here, the price is 1 minus the total quantity on the market. The total quantity of the market is what's produced by everyone else, which is q times n minus 1. Because there's n minus 1 other firms. They're each producing q.

If I choose q prime, then this gives me the total amount produced. This gives me the price. And then I subtract my cost. And if we could actually do a little algebra. I don't expect people to see this, but I think it's actually going to be this. The formula doesn't really matter, but this is what it will be. But we'll just stick with these two formulas here.

And then what is our strategy profile? Well, it's exactly what we said. So one, each firm produces qM. Until a deviation. But two, if anyone has deviated. every firm produces qNE So this is just formally writing down the Nash reversion idea that we said.

We start by producing qM every period. As long as no one's deviated, that's what we do. But if any firm deviates, then we all revert to the Nash equilibrium quantities, qNE.

So now we want to check that this is a subgame perfect Nash equilibrium. As usual, we're going to use the one-shot deviation principle. And we're going to check every history and check that no firm has a profitable one-shot deviation. But we're going to use a different representation of this strategy called an automaton, which is going to turn out to be very useful for structuring the way we think about these equilibria. So let's look at what's called the automaton representation-- automaton representation.

So you recall that when we describe these strategies, the way we analyze them is we often grouped all the histories into two buckets. We said they're the histories where someone's deviated, and the histories where no one's deviated. And the automaton representation formalizes that in terms of states.

We think of, what's the current state that we're in, and how do we behave in that state? So let me just write it down. To represent this equilibrium, we're going to need two states. We'll call it the monopoly state and the stage game Nash equilibrium state. And now we need to specify how we behave in each state.

Well, in the monopoly state, we each produce the monopoly price-- monopoly quantity. Sorry. That's easy. So qi equals qM for all i. And in the Nash equilibrium state, we each produce the Nash equilibrium quantity. That's pretty easy.

So what do I really mean by these states? Well, they're this mental construct. But let's just describe how things work, and then we can maybe interpret it.

So the idea is we start in this state. Initially, we're each going to produce qM. So I'm going to put an arrow here to say this is the state that we start in. And in this state, what we're supposed to do is we're each supposed to produce the monopoly quantity.

Now what we have to describe is how we transition between states. So the idea of this equilibrium is we're going to stay in this state as long as no one has deviated. But if someone deviates, then we move to a different state.

So what I'll write is I'll write an arrow here to say, we're in state m. And we stay in state m if what? If qi equals qM for all i. That's the only way we're going to stay in this state. Otherwise, someone has deviated, and then we move to the other state. So I'll write otherwise here.

If we're in this state, well, once someone's deviated and we're in this state, we actually stay in this state forever. There's no way to leave this state. So we're just going to put an arrow here that says, anything.

So this simple graph actually fully describes a strategy. And, in fact, fully describes the strategy profile in this equilibrium. And I think the way I like to visualize it-- I mean, of course, if firms are actually colluding today, they'd use computers. They'd keep track of things.

But a very simple way to think about this is all I need to do is have two positions on my desk-- the monopoly position and the Nash equilibrium position. And I like to imagine I just have a stone. And I start by putting the stone in the monopoly position.

I look at my desk. I see-- and every firm is doing the same thing. We start with the monopoly position. We all say, what do we do here? We all produce the monopoly level.

Then tomorrow we observe what everyone produced yesterday. And then we potentially move the stone from the monopoly position to the Nash equilibrium position if we observe that someone has deviated. If no one deviates, we keep the stone there. And then tomorrow we see where the stone is, and we decide how to play. So this is a very simple representation of the strategy profile.

So now we come over here, and let's write down the conditions for equilibrium in this game. So it turns out the key to analyzing equilibrium using this automaton representation is to write down what we'll call the value function. And in this case, we need to specify the value function each of the two states. So we're going to have vM and vNE.

And what the value function says is, in the first case, suppose we're in state M today, and everyone follows this strategy. What is the average discounted payoff of the firms going to be? And then vNE says, well, if we're in state NE, and we all follow that strategy, what is the average discounted payoff of each firm going to be?

So let's start with M. If we're in state M, It's a symmetric equilibrium, so we're just going to write down one firm's payoffs. All the other firms are going to be the same. If we're in state M, and we all follow the strategy, what is the expected discounted average payoff of each firm?

Well, first let's see what happens. If we're in M, well today we all produced a monopoly quantity. We say in M. Tomorrow we all produce the monopoly quantity. We stay an M. The next day, we all produce the monopoly quantity. So we're going to get the monopoly quantity every single period.

And what is the payoff from that-- the average discounted payoff from that? Well, it's just f of qM. f Of q is a function that tells me what is my payoff if we all produce quantity q? So if we all produce quantity qM, this is exactly my payoff.

What about vNE? Well, here it's going to be f qNE. This is what are my average discounted payoff from here on if we're in state in this state. So now we're going to use the value functions to write down the equilibrium conditions. And when we use an automaton representation, we have one equilibrium condition for every state.

So we're going to have state M, and we have state NE. And what our equilibrium condition says is if we're in state in this-- let's take this condition. If we're in state M, then each firm is better off continuing with their equilibrium strategy than choosing the best possible one-shot deviation.

So each equilibrium is going to have the following form. In state M, we need to check that if I do what I'm supposed to, well, by definition I get my payoff vM. That's what it means if we continue with the equilibrium strategy in state M. And we have to compare that to the best possible payoff I could get from a one-shot deviation.

And then we'll do the same in NE. Here we are today. This is what I'll get if I follow the equilibrium strategy. And on the right-hand side we want to say, what is the best possible payoff I could get from a one-shot deviation?

So let's start-- I think maybe this one is easier. Oh, no. We can start with M. It's fine. So what is the best possible-- if I'm going to deviate, how should I deviate when we're in state M. Yeah?

STUDENT: [INAUDIBLE].

IAN BALL: Exactly right. So if I'm going to deviate at all, I know I'm going to be punished. I'm going to be punished in exactly the same way no matter how I deviate. So I might as well choose the deviation that gives me the highest payoff today. And that's exactly-- and my payoff from that is exactly given by this function g.

So the best possible deviation-- well, today I'm going to get g of qM. Remember, this doesn't mean I'm playing qM. This means if everyone else is playing qM, and I choose the quantity q prime that maximizes my flow payoff, this is the payoff I'm going to get. And then if I do that, what state are we going to be in tomorrow?

Well, we're in this state. I deviate, which means we move to the Nash equilibrium state. So from tomorrow onward, we're going to be in the Nash equilibrium state. So my payoff is going to be vNE. But the only thing we're missing is discounting.

Remember, with average discounted payoffs, how much weight do I put on what I get today, and how much weight do I put on the infinite future? Well, we know we always put 1 minus delta today and delta in the future. So we're going to have 1 minus delta here plus delta vNE. And it's really important that we interpret this correctly.

The first term is what's the payoff I get today if I choose this optimal deviation? This term is not saying, what's my payoff tomorrow? This is incorporating my payoff tomorrow, the next day, the next day, and the next day.

But by our definition of vNE, we've defined vNE to say, what is my payoff from today onwards once I'm in state NE? And that's exactly what we're going to measure on this side. So this is reflect-- it's a little tricky because in this case, it happens to agree with my flow payoff. But it's just important to understand that this is representing everything from tomorrow onward. It's not just representing what happens tomorrow. Because what I do today affects everything from tomorrow onwards.

And now let's go to vNE. We're going to have a similar formula. What's our formula going to be down here? Well, I have the same optimal deviation. If I'm going to deviate, I might as well do the best thing. So it's a little tricky in this case. Yeah. So maybe--

Here, it's actually a little tricky because my best deviation since we're in a Nash equilibrium is actually to do what I'm supposed to do. But it actually doesn't matter here because we're going to stay in NE anyway. So let's see. Maybe this case, we should just-- yeah, I feel like this equation is misleading.

It's true, but let's just say-- I'll just say best deviation, but my claim is that this is never profitable because if we're in the Nash equilibrium state today, whatever we do today, it can't be any better than playing the Nash equilibrium action, and it has no effect on what happens tomorrow. Yeah?

STUDENT: I'm a little confused. You said how you define the value function, it's just equal to f. And as I understand, f is just my payoff today.

IAN BALL: Ah. Great point. That is very confusing. Yeah, I see where you're going. It happens to agree with f, but that's simply because we're using average discounted payoffs. So maybe what I should write is 1 minus delta times f of qM-- let me get at good point-- plus delta f of qM, plus dot, dot, dot. That's really what I mean. Good point.

So if I want to compute what happens, I'm starting in M. I'm going to get f of qM today. I'm going to get f of qM tomorrow and so on. I'm going to discount, and I'm going to take the average. It turns out that equals f of qM by the way we've defined it.

So you're right. It's tricky in this case because it happens to reduce to this. But this is the interpretation of where that's coming from. And that reduces to this. Is that clear? Yeah, great. Thanks. I should have written it that way the first time. And of course, the same applies here. Yes?

STUDENT: For the Nash equilibrium, you just said vNE is [INAUDIBLE] [? never ?] profitable. Is that just by definition that we are in a Nash equilibrium?

IAN BALL: It's a little trickier than that. So I'm hesitant to write. Maybe I should go back and write it the other way. What's tricky is it's the fact that we're in a Nash equilibrium and that tomorrow we're going to be in the same state, no matter what. So maybe let me write it out more formally.

So what we need to show-- so in the first case, I can use this function g to simplify my life. Here, it's a little trickier. So really what I need is it's greater than pi qi prime. Yeah, let me write it this way. And actually, the first one we could write this way, but it reduces.

So I'm firm i. What I could do instead of following my equilibrium strategy is I could instead produce qi prime today. If I produce qi prime today, and everyone else is producing the Nash equilibrium amount, then this is my flow payoff today. And we multiply that by 1 minus delta.

Whatever I do today, we're going to stay in the state Nash equilibrium. Because if we look at the arrows here, whatever I do, the arrow brings us back to the Nash equilibrium state. So we're going to get delta and Nash equilibrium tomorrow. And how do we see that this is less than or equal to this?

Well, whatever qi prime might choose, it can't give me a higher flow payoff than choosing the Nash equilibrium quantity because that's the definition of a best response. So we see that this is less than or equal to 1 minus delta f of qNE plus delta vNE. And that's exactly vNE.

So how do we see that this right-hand side is smaller than what we want? Well, all I've done is I've compared the flow payoff to the Nash equilibrium payoff. Whatever quantity I choose, it can't be as good. It's no better than playing my best response, which is exactly the Nash equilibrium quantity, and that gives me these payoffs. And then I substitute in this over here, and we're going to get this. So any questions on that?

So the action is going to be here. We want to check that star is satisfied, and let's plug things in and see what we get. So let's look at star.

Well, we have our inequality, star. And we have our values for the v's on the left side. So we're just going to substitute these values in to our inequality over here and see what we get. So vM becomes f of qM. And that needs to be greater than or equal to 1 minus delta g of qM plus delta times vNE, which we look over here, that's f of qM.

And here, I think we can really capture exactly the idea of the temptation of the deviation. If I don't deviate, I get f of qM forever after. If I do deviate, I get maybe a higher number today because I increase my flow payoff today. But then forever after, I get a lower number.

So let's note that g of qM is strictly greater than f of qM, which is, in turn, strictly greater than f of qNE. So this is what I'm supposed to get. If I deviate, I bump up my payoff to this today, but then I drop down to this forever after. And we want to make sure that I'm patient enough that the dropdown forever after offsets, or sufficiently offsets the potential game today. Any questions?

So now let's just do some algebra. Let's bring-- we want to bring everything to this side I think. So what we're going to have is delta-- I want to bring everything involving delta to the left side, and everything not involving delta to the right side.

So I have minus delta g of qM. I'll bring that to the left side, and that becomes plus. Then I have plus delta f of qNE. I'll subtract that and bring it to this side. So I get minus f of qNE greater than or equal to--

So now let me just divide, and I just get delta greater than or equal to this threshold, maybe delta upper bar, which is defined to equal g of qM minus f of qM. And I sometimes think it's easier to decompose it this way. This is g of--

So as we expected, this is going to be a subgame perfect Nash equilibrium if delta is large enough, if the players are patient enough. And let's try to interpret the payoffs over here.

Well, what I want to think is I want to think that this gap represents my one-shot gain from deviating. But this gap-- this is my gain today from deviating, And this is my loss forever after. Maybe we'll call this g equals gain, and this loss equals l.

So we're supposed to be playing the monopoly quantity every period. If I deviate, I can overproduce. I do strictly better today, I gain g for myself. But forever after, instead of getting the monopoly quantity, we all produce the Nash equilibrium quantity, which is strictly worse. So we lose this amount l.

And it turns out, if you just look at the algebra I've done here, g of qM minus f of qM, that's exactly the gain. So we get g over here. And at the bottom, we exactly have the gain plus the loss. And this is kind of a general formula that says, what's the requirement on delta? It has to be large enough, roughly, that the potential one period gain I get is offset by the eventual loss I experience.

So let's make sure we understand how this varies with g and l. So let me ask this. So if g gets bigger, what happens to delta bar? Let's start with L. That's the easier one. If L gets bigger what happens? Yeah?

STUDENT: I'm going to be honest. I'm having trouble working out the math.

IAN BALL: But yeah. So this is just algebra. You can check that on your own. Let's just look at this and this formula over here, just g over g plus l.

STUDENT: As the loss gets bigger, delta is going to have to be smaller.

IAN BALL: Delta bar gets smaller. And what is that? So can we interpret that? As the loss gets bigger, delta bar gets smaller. So what's the--

STUDENT: That just means you have to put more value-- or put less value on the future if you're going to be moving more in the future.

IAN BALL: Right. So it's easier to sustain equilibrium. If the loss gets larger, that means the punishment is very powerful. If I get caught for deviating, tomorrow I experience this huge loss. If I experience a huge loss for deviating, it's not very attractive for me to deviate. Therefore, it's easier to sustain the equilibrium.

Because if delta bar is smaller, that means I don't have to be as patient for this to be an equilibrium. If l is really, really big, and the punishment is really, really bad, even if I'm pretty impatient, this still constitutes an equilibrium because the punishment is so severe.

So more delta satisfy this-- even smaller deltas, more impatient deltas, when l is very large. What about g? Can we also reason through g? This is a little trickier because g is on the top and the bottom, but it is increasing in g. Maybe it's not quite so easy to see, but it is. So why is this increasing in g? Yeah?

STUDENT: Well, if you're able to gain more than the threshold for deciding that you shouldn't deviate would need to be higher.

IAN BALL: Right. If I can gain a lot today, it's going to be really, really tempting for me to deviate. So it's going to be harder to sustain the equilibrium. And that's only going to work if the players are really, really, really patient. So this threshold gets higher, meaning fewer deltas satisfy it because the temptation to deviate gets really, really big.

And you can check that-- everything's in terms of f and g. But remember, at the end of the day, f and g involve n. So let's just see if we can get some intuition for this. Well, maybe this is too hard to see, but it turns out that delta bar, you can check, is increasing in n.

So can we give an economic story? What does this tell us? Or what might we expect about which cartels are successful and which cartels are not successful based on this fact?

STUDENT: Smaller cartels.

IAN BALL: Smaller cartels meaning what? Yeah, you're right. But--

STUDENT: Like the ones that have less players involved.

IAN BALL: Fewer firms, right? The players themselves are big. So we have to be careful about what small means. The players are bigger, but the number of players is smaller. So as n gets bigger, we have to be really, really patient to sustain cartel behavior. And given that people aren't necessarily that patient, that may not work out.

So we would expect that cartels might be more successful and last longer when n is small and we just have a few big players interacting. And indeed, that's pretty consistent with the empirical evidence, that it's hard for many, many firms to coordinate and work together and sustain cartel behavior. Even when there's few firms, things do break down.

So this may seem kind of just a math story. We have the Nash equilibrium. We have a monopoly. But this is actually exactly what happened in the 1980s. Saudi Arabia got frustrated that everyone else was overproducing a little bit, and they flooded the market, and the price of oil dropped dramatically.

So we see exactly this kind of punishment. Saudi Arabia said, look, I'm tired of these other countries overproducing relative to what they're supposed to. And they punished everyone by flooding the market, and the price dropped. And this had huge effects globally.

Let's do one more equilibrium construction that I think just kind of. Illustrates things. So let's do a final equilibrium. So before, we did a Nash reversion equilibrium. Now we're going to do what's called a carrot and stick equilibrium.

And what was distinctive about this equilibrium is once someone deviated, we reverted to Nash equilibrium forever after. But if you look at what happened with Saudi Arabia, they didn't just produce a high amount of oil forever. They flooded the market for a few years, and then they eventually brought their production down again, and prices went up again.

So what we often see in these cartels is punishments of limited duration. Someone overproduces. We punish them for a year or two, or a month or two, or whatever. And then we get back to the good thing that we are doing. Because ultimately, that's what we want to do.

So the question is, can we have these carrot and stick strategies? And the way this is going to work is we're going to have two production levels-- qc less than qs. So carrot and stick are generic terms for carrot is a reward and stick is a punishment. This is like, I guess, animals or pets. You give it a carrot or you-- I mean, it's kind of violent. But anyway, that's what it is.

So OK. I'll use this-- maybe I shouldn't use this terminology, but I've inherited this terminology. I'll use it. So the carrot is the low production level. This is the good thing that we like. The stick is the high production level. So this could be the Nash equilibrium and the monopoly, but it doesn't have to be. It's just general production levels.

And the idea is we're going to start with the carrot. And if someone deviates, we're all going to produce the stick level tomorrow, but only for one period. Then if everyone does what they're supposed to, we come back to the carrot level. So let's try to represent this as an automaton. So let me erase this automaton and get a new one.

So now our automaton is going to have two states-- the carrot state and the stick state, so c and s. And we're going to start in the carrot state. And in the carrot state, everyone produces this low-level c, which is good. And in the stick state everyone produces this high-level qs, which is not so good for everyone.

And now we have to say what happens. I think in this case, the transitions are a little more complicated. So in the carrot state, we're going to stay in the carrot state if everyone produces the carrot level. So we have an arrow here if what? If qi equals qc for all i. If everyone does what they're supposed to, we stay here. Otherwise, we go to the sixth state.

But now this is where it gets a little tricky. Let's see. Can we figure out what happens in the stick state? What happens if we're in the stick state, and everyone produces the stick quantity? What state do we move to? Do we stay, or do we move?

STUDENT: Move to the carrot.

IAN BALL: We move back. Because now if everyone does what they're supposed to-- OK, we did our punishment. We're now ready to get back to where we were before. So now we're going to have this arrow if qi equals qs for all i. And then what happens if someone doesn't play the stick? What do we do? Well, then we stay in the stick state and punish them again tomorrow.

So at first this may look confusing. But I think the way to interpret it is in either state, if everyone does what they're supposed to, we move to the carrot state tomorrow. That's this arrow and this arrow. If anyone does what they're not supposed to, we go to the stick state tomorrow. That's this arrow and this arrow. And that's true whether we're in the carrot state today or the stick state today. Either way, we go to the stick state tomorrow.

I should say when I write "otherwise," it can be a bit ambiguous. Otherwise is referring to the source of the node. So in the c state think of, we're in this state. Look at all the arrows. We could either do this, or otherwise we do this.

And then similarly, if I'm in the stick state, if we do this, we go here. And otherwise, we go here. But when you just see all the otherwises you might think, what are they referring to? They're referring to the state we're in, the source of all the arrows.

So now, I don't want to do tons of algebra, but I just want to see if we can set up the value function and the equilibrium conditions for this equilibrium. And then on the problem set, you'll get to actually do all the algebra. But I think here, we'll see the benefits of this value function approach.

So there's two steps always. On the left side, we compute the value functions, as we did above. And the right side, we write out the equilibrium conditions. If one of these is the Nash equilibrium level, it's a little tricky.

So maybe just to simplify my life, let's assume that these are not equal to qNE. So the carrot and the stick are both different from the Nash equilibrium level, just it avoids this kind of issue we got before, where the best deviation is not a deviation. And let's just not worry about that.

Now, I want us to conceptually understand what's happening here. The left-hand side is nothing to do with equilibrium. The left-hand side is just saying, if this is the strategy we were to play, what would everyone's payoffs be in each state? Then using that information, we go to the right side, which is saying, is it actually an equilibrium for people to play this way?

But notice, we can ask, what would my payoff be in each state if this is the strategy we play? I can ask that question whether or not that strategy is actually in equilibrium. So let's try to write down our value function formulas here.

So we have Vc and Vs. So we have a carrot state and a stick state. And we want to compute. Remember, the interpretation is the average discounted payoff from today onward if we're in this state.

So let's say we're in the c state. What is our payoff going to be today. If we all follow the c state? We can use the same notation that was up here? Well, it's f of qc. And we'll multiply that by 1 minus delta.

And then if we follow f of qc, what happens tomorrow? What state are we in? We're still in the carrot state. So now we get delta v of c.

Now, this is the one that's a bit trickier. If we're in the stick state, today we get 1 minus delta f of qs. But then what do we get on the right-hand side here? Yeah?

STUDENT: Is it delta qc?

IAN BALL: Exactly. Because if we do what we're supposed to do today, we end up in the carrot state the following period. This is saying, what is my average discounted payout from today onward in the stick state? Well, it's the payoff I get today plus delta-- we've got a discount-- times my average discounted payoff from tomorrow onwards. Yes?

STUDENT: So value function is assuming that you follow the strategy profile.

IAN BALL: Exactly. So that's the key thing over here. We're not saying about equilibrium. We're not talking about deviations. We're just saying given the strategy profile, let's compute the value functions in each state.

Now in general, you can see we're going to have a system of equations with however many states there are. Call it n. No, that's the number of players. t-- k. If there's k states-- so I'll write this down here. If there's k states, in general, we have k equations and k unknowns because we have to specify the value in each state, and we have a formula for each of these. So we get k states and k unknowns.

And, as you know, this is going to be a linear system. And it turns out it will actually always have a solution. That's something that's a math result you can show. So this is always a way that we can compute what the values are going to be in each state.

Notice that this is kind of-- there's something very clever going on. Because if I just started in one state-- let's suppose I didn't use this approach, and I just wanted to say-- let's say we had a really complicated-- let's say we had 10 states. And I wanted to say, what is my payoff starting in this state?

Well, I'd have to keep going through it. I'd say, well, I do this today, then I go to this state. Then I do this tomorrow, I go to this state. I go to this state. I go to this state. You might think I'd have to compute the whole infinite horizon of all the different states I could pass through.

What's kind of magical about this approach is we've just reduced all these computations to a single system of linear equations, which is much easier to solve. I don't have to explicitly compute my entire path through the automaton. I just have to solve the system of equations. And this is related to some kind of nice math. Yeah?

STUDENT: Each unknown corresponds [INAUDIBLE] one-to-one correspondence between states.

IAN BALL: Exactly. So formally, you can think of v as a value function that's defined on the domain of the set of states. So a function needs to specify a number for every state, and that's precisely the value we get starting in that state.

And if we solve this-- we can pretty easily solve this first equation just to make sure this makes sense. If we subtract delta Vc from the right side, then we get 1 minus delta Vc equals 1 minus delta f of qc. So that tells us that Vc equals f of qc. Does that seem right?

Yeah. Because if we're in state c, we never leave. So every period we get f of qc, and that's our average discounted payoff. Question? Yeah?

STUDENT: I mean, because we know we stayed there-- or for Vc, we know we stay in state c if we do it the first time. Would there have been [INAUDIBLE]. Like, mathematically, you would get the same answer [INAUDIBLE] instead of writing [INAUDIBLE].

IAN BALL: You're right. You could have just jumped-- you're saying you can basically see what the payoff is going to be here.

STUDENT: Yeah. But delta times f of Vc.

IAN BALL: You could have, yeah. But it's down here where the benefit comes in. Because here, and more generally, what happens tomorrow depends on a third state. So this is the general approach that always works.

But you're right. In this case, if we're in what's called an absorbing state, once we're there, we never leave, it's pretty easy to calculate what happens. We don't need to worry about what happens in any other state. But when we have transitions between the states, then it becomes trickier. But yeah, exactly right. But this will never give you the wrong answer. Yeah.

OK. So now let's write down our equilibrium conditions. Once again, we have a condition for each state. But remember, we're actually capturing a lot here. Because in each state, there are a lot of different one-shot deviations we could use. So when we say equilibrium condition, what we really mean is we're saying in each state we want to check that no one-shot deviation is profitable.

And the way we can check that is to check that the best one-shot deviation is not profitable. Because if the best deviation is not profitable, then none of the deviations is profitable. There may be some really bad deviations. We don't care about those. We only care about the most attractive deviations. And that's exactly what we're going to put in here.

So in the carrot state, we need Vc greater than or equal to something, and Vs greater than or equal to something. Eventually, we'll substitute things in. But I want to just do it conceptually this way.

So in the carrot state, once again, if we deviate we're going to get punished. And the punishment we experience doesn't depend on the way in which we deviate. So we might as well choose the deviation that maximizes our flow payoff today. And what is our flow payoff going to be if we choose the deviation that maximizes our flow payoff? We have a function for this. It's g of qc.

So this says, if I choose the best deviation today, this is what I get today. If I deviate, what state are we going to be in tomorrow? Yeah?

STUDENT: Vc.

IAN BALL: Yeah, exactly. So now we're going to get the s over here. And now I think this is the hardest one, and we'll stop here. Now we're in the stick state. We're going to choose our best deviation in the stick state.

Again, I'm only allowed to write it this way because the punishment I experience doesn't depend on exactly which quantity I deviate to. If the deviation I chose affected the punishment, then we'd have to-- things would be a lot more complicated.

But all we have to ask is, am I deviating or not? However I deviate, it's treated the same way in the future. So I might as well choose the deviation that maximizes my flow payoff. And then what do we get on the right-hand side? Delta times what? What happens if I deviate? Vs again.

So notice when we were defining these, we always had Vc here because we were looking at what happens if I do what I'm supposed to do. And therefore, we go to the carrot state. When we're looking at the equilibrium condition, we're considering deviations on the right-hand side. Whenever I deviate, regardless of which state I'm in, I go to the stick state.

Now, let's just understand what I think is kind of tricky here, is this is actually capturing a lot. If I deviate today-- so let's just understand the anatomy of my one-shot deviation. So we have today, tomorrow, and then thereafter.

If I deviate today to, say, qi prime in the stick state, we had a really simple formula on the right-hand side. But what's actually happening is pretty complicated. If I deviate today to qi prime and then return to my equilibrium strategy, what am I going to produce tomorrow? [INAUDIBLE] want to guess? You see it?

Well, tomorrow I go to the stick state. So I'm going to produce qs tomorrow. And then the day after that, I'm actually going to produce thereafter, I'm going to produce qc.

So this one-shot deviation is actually really complicated. I'm in the stick state. I deviate. That means I stay in the stick state. Tomorrow, I'm following my equilibrium strategy in the stick state, so I choose qs. But if I choose qs and I'm in the stick state, then we transition to the carrot state the next day. And then I choose qc the next day.

So the path of play following this deviation-- and then it's going to be qc forever after-- is actually really complicated. When we wrote our formula here, things looked much simpler. How is this formula reflecting this complicated path moving forward? Yeah?

STUDENT: I just have a question. Why are sometimes s And c subscripts, other times superscripts? Good point. I think I want these to be superscripts.

IAN BALL: There's not a great reason. I often had q sub i here, so I made this a superscript. But then I never have v sub i, so I made this a subscript. But everything's the same. Yeah?

STUDENT: Also, was qi prime-- that could be an either stick or a carrot state, right?

IAN BALL: I'm looking at this-- I'm looking specifically at this case. I'm in the stick state. I deviate. Because I deviate, I'm in the stick state tomorrow. So then I follow my strategy since it's a one-shot deviation and choose the stick quantity. But then the next period, I'm in the carrot state. And then I choose qc.

So my point here is that a one-shot deviation can lead to a really complicated path of play. The reason I don't have to worry about it is Vs is already taking into account this future path. So this path here of SCCC is already reflected in the value function Vs.

In fact, we can see over here Vs is precisely reflecting that I get the stick today and then the carrot forever after. So this is showing you the power of the value functions, that I don't really have to keep track of the total sequence of play. I just have to keep track of what state I get to tomorrow, and then it takes care of itself.

Let me stop there. Sorry for going slightly over. And good luck on the midterm on Tuesday.