[SQUEAKING]

[RUSTLING]

[CLICKING]

GLENN ELLISON: OK, so I'm going to start with something I also cover in [14.]122, which is collusion with perfect, simplest model of dynamic collusion. So we say the, I guess, reason for thinking about dynamic competition is, in part, because we think dynamic competition is a very good model for how many industries work. It's T-Mobile and Verizon and whatever competing against each other period after period after period. It's the same cable companies competing. It's the same companies competing, and many industries with a small number of firms.

The small number of firms reflects that. It's hard to enter, and you have those same firms for over and over and over again. And then the second reason we're covering this is a lot of this course is about markups and thinking about where markups come from. I've talked about differentiation as a source of markups, and I've talked about search as a source of markups. I think we also think that dynamic collusion is another potential source of markups. And I did do the--

Ray Miller and Weinberg had that conclusion that it looks like, when you look at the Miller-Coors merger, part of the effect was Miller and Coors no longer pricing in static Nash form as they had before and AB InBev as they had before the merger. The theory of dynamic collusion is thinking about how it is that we can think about when firms are going to have markups or when they're going to do something other than static Nash. OK, so simplest model. I've got N identical firms. They have a constant marginal cost of c. They all have a discount factor of delta.

We're going to have them play the simplest Bertrand pricing game where there's some market demand function, Q of p. So that means all of the demand goes to whichever firm charges the lowest price. The other firms sell 0. If firms charge equal prices, I'll assume they split the demand. Firms all observe prices at the end of every period. So at the history at time t is h t minus 1. It's all the prices charged by all the firms in every period, tau, less than t minus 1.

And a strategy in a repeated game is a function. Firm i sets its time t price as a function of everything that it's observed in all previous periods. And when I teach this model in [14.]122, I note that this model has a subgame perfect equilibrium that looks like perfect competition, the static Nash equilibrium, which is just everybody plays price equal to cost in every period. But if delta is at least 1 minus 1 over N, this model also has a subgame perfect equilibrium that looks like perfect collusion.

And what that is is the firms charge the monopoly price as long as everyone has charged the monopoly price in every previous period. And if anyone ever deviates, they set price equal to c in all future periods. That price war punishment, as we would call it-- that price war punishment is going to deter people from deviating and lead everyone to want to coordinate a monopoly price, if they're patient, and this is how patient they have to be. So some comments. This model identifies another factor that can determine markups.

In particular, just repeated competition can lead to markups with no product differentiation, no search costs, nothing else. It has a lot of natural comparative statics. So it's easier to sustain collusion when delta is larger. So when delta is above this number, you can collude. When delta is below that number, you cannot. So what could delta mean? If you think about delta in this model, delta is the discount you apply.

In between, you see someone cheat on the agreement, and then you start punishing them. And so delta is your time rate of discounting between those two events. Someone cheats, I start punishing. So what makes collusion harder? One would be detection lags. If people cheat on the agreement, you don't see them cheat for a while. That makes it harder to collude.

Lumpy demands. In some industries, firms are selling goods over and over and over again and see their demand coming in real time. Other demand, like if you think about Raytheon competing with Lockheed, if they're still separate companies, to produce next-generation fighter jets for the US government, if those orders for new-generation fighter jet come along once every 15 years, if someone cheats on the agreement to collude on prices and you can't do anything about it for 15 years, delta is going to be very small. So whether demand is coming in continuously or lumpy matters.

It's easier to collude in growing markets. It's easier to collude in growing markets because, again, the basic consideration here is how much do you hate-- how much do you dislike this punishment of switching to pricing at marginal cost and losing the collusive profits? And how much does that deter you from deviating? If the future is more important than the present, then that's like delta is large. That makes collusion easy. And on the other hand, technological change, if this is a market that's going to be replaced by some other product, then the future doesn't matter.

If the future doesn't matter, then it's hard to collude. And so you might see collusion breaking down just because, when the market is on its way out, collusion could break down. Model also has a nice comparative static, that it's easier to sustain collusion when N is small. Because, again, when N is small, 1 minus 1 over N is a smaller number, so more markets where you could collude. Another common on this is that this theory also, in some ways, suggests that everything we've done in the entire semester is almost worthless.

Because anytime you have a market with a small number of firms, how much differentiation there is, how much search costs there are, how much of whatever, how much multiproduct competition, all of that stuff, in some sense, is irrelevant because the firms could collude on anything. And so maybe one suggestion is that we either have to think about which equilibrium in this repeated game we think is going to happen, or we have to just accept that we can't always forecast what's going on and we're always-- we may always be surprised about how things are happening.

The other thing I'd like to say about this model is that, numerically, it's just giving the wrong answer. It's saying that, with delta at least 1 minus 1 over N, we can collude. And if you think about it, if you think about a typical industry-- prices are posted on the internet or whatever. How long does it take you to observe that your rival has colluded? Like, if you're a T-Mobile and you see Verizon changing their prices-- they put a big TV ad campaign saying they've changed their prices-- within a week or within a day, you know what they've done.

What's the discount factor over a day? Discount factor of a day is 0.000. Or delta should be 0.9999-- or five 9s or whatever. You could collude with thousands of firms. And so I think that, while-- even you put N equals 100, delta is 0.99, 0.99 is a good discount factor for a month or less than a month. So this would say, with 100 firms, even if it takes you a month to realize your rivals have changed their prices, you could collude perfectly.

And I think we don't think that collusion typically happens with 100 firms when you have month-long delays. So there's something about collusion is just harder than this model suggests it should be. And so one thing you find in the IO literature is a bunch of factors that all may make collusion harder than it would be in this model. So I thought I would do is I'm just going to-- it's a fairly old literature, and it's all well-covered in Tirole's book. So I'll do a few of those models, but maybe I'll go through some of them kind of quickly.

So the one I did want to do in more detail was this literature on collusion with imperfect monitoring. And I think this is, I think, what we think of as probably the most important reason why collusion is hard, which is that it's often hard to know when your rivals have cheated. It's easy to know when your rivals are cheated if your rivals are consumer-facing businesses that have to post prices that are visible. But in many industries where you're selling to businesses or selling elsewhere or selling with negotiated prices, it can be much harder to know if your rivals are cheating on an agreement.

So I'm going to discuss this as-- ideas due to Green and Porter. This version of the model I'm giving you here is the one that's just out of Tirole's textbook, where Tirole reconfigured it to be an easier model to solve but, obviously, not as realistic. But so, in this story, I've got two firms. They have a constant marginal cost of c, and they're engaged in Bertrand competition. And market demand is noisy.

So market demand is either Q of p, with some probability, 1 minus alpha, or it's 0 with probability alpha. So where that might make sense is you could have two firms that are competing to supply computer equipment to some large business. And many months, that business calls them up and says, I need 132 new computers, or I need 147 new desk chairs, or whatever. But then, some months, it doesn't call.

So there's these two levels of demand. Either there's demand Q of p because the firm places a large order with one of you and it's 0. I'm going to assume that demand goes to the lower-priced firm or splits 50/50 if their prices are equal. And I'm going to assume there's imperfect observability in that firms see whether they've made a delivery to the customer. But they don't know what their rival's price is or what the rival's sales are. So you both have websites.

The websites give your official price. But you don't know whether the firm called your rival and said, can you give me a 10% discount and I'll buy all my stuff from you this month? And then they did. And then they shipped it to them. All you can observe is your own demand. So the difficulty when that happens is we get this imperfect observability. Suppose there's some month where I get zero demand. One of two things happened.

One thing that could have happened is we're just in a low-demand month and the firm didn't place any orders. The other thing that could have happened is my rival is cheating on the agreement, and the firm did make an order, and the entire order went to my rival. And obviously, you could call your rival and say, oh, I got zero demand. How about you? And then they can be like, yeah, zero demand, me too, or whatever. And you just don't know. Are they lying to you, or is this really what happened?

So immediate conclusion-- no matter what the discount factor is, this model can't have an equilibrium in which both firms always charge the monopoly price on the equilibrium path. Because if you had an equilibrium where both firms were always charging the monopoly price, there would be some chance, given that these are i.i.d., that you'd have a million periods in a row where there's no demand. And if it was in every period on the equilibrium path, you always charge a monopoly price, then after that million periods of million zeros in a row, you still have to have both firms charge the monopoly price.

But then if that was going to happen, obviously, it's better for you to cheat in the first million periods in a row and sell the entire market instead of half the market and leave your rival with nothing than to follow the equilibrium. So the fact that we've got that fundamental unobserved ability means there's no-- perfect collusion is just impossible, no matter how patient firms are. And then I think the key insight of Green and Porter is that while perfect collusion is impossible, you can get partial collusion in a model like this.

So proposition, if alpha is less than a half, which you should think of the noise is not too large and delta sufficiently close to 1, the model has a partially collusive equilibrium. And what happens in this partially collusive equilibrium is the firms initially set the monopoly price. And then if they ever get zero demand, they switch to c for capital T periods and then they go back to the monopoly price. So what we would see-- what we would observe in this model is we would observe the firms going along, charging monopoly price.

They would do that for some period of time. And then when there's one of these bad demand periods where the customer doesn't order from them, they would switch, and they would charge c for capital T periods. And then price would go back up. And then they would continue to collude for a while. Another price war. And so, in this equilibrium, we'd get this sort of funny price pattern that looks like that, where the firms go along, charging monopoly price. They jump down. They have a price war. And then they go on like that and so on.

So then the theorem is that that is an equilibrium of this model. And intuitively, why don't you want to deviate from such an equilibrium? If you cheat in this model-- the other firm's charging Pm. You charge Pm minus epsilon to try to get all the demand. You do get some extra profit today from getting the whole market instead of half the market, but then you immediately trigger this T period price war.

The money you lose in these T periods of a price war is going to offset the gain you get in the one period in which you succeeded. And T sometimes doesn't need to be bigger than two or three to make that true. OK. So how am I going to prove that? The way I'm going do it is I'm just going to do-- there's lights on the board. [VOCALIZING]

OK, the way I'm going to do it is it's just going to be a simple dynamic programming kind of calculation. I make these two definitions. Vm is the present discounted value of payoffs at t equals 0. And Vp is the present discounted value of payoffs if we're at the start of the punishment phase. And the strategies are symmetric, so these things don't need an index for the players on them. And I just need to show that there's no profitable deviation.

Showing that there's no profitable deviation during the punishment phase is easy. If we're down here and I decide that I'm supposed to charge c and I'm going to charge some other price, well, if I deviate and I charge a price that's less than c, all I do is sell at a loss, and I don't affect the future at all. And if I deviate and raise my price, I sell nothing, which is the same. And again, I don't affect the future, as I've described the strategies. So in the punishment phase, you're never going to want to deviate.

You're never going to gain from deviating because all you do is lose or weakly lose in the current period and have no effect on the future. So it just remains to show that when you're at some point, like this, then you don't, also, want to deviate from the equilibrium. So to do that, what I need to show is, if I follow the equilibrium, I'm going to get this present discounted value Vm. If I do deviate, if there is-- we're in one of the positive demand periods. I'm going to get the full monopoly profit instead of one half the monopoly profit.

But then, one period later, we're going to be starting a punishment phase. So I'm going to get delta times Vp. So I need to show that Vm is at least the right-hand side of that expression. So to show that, I just need to find Vm and Vp. And to do that, I just solve two equations and two unknowns. One recursive equation for Vm is the profit I get at the start of the game is with probability 1 minus alpha. I get pi m over 2. And then next period, we remain in the cooperative phase.

And with probability alpha, one period later, we end up in the punishment phase. So this is one equation for Vm is 1 minus alpha times pi m over 2 plus delta Vm plus alpha delta Vp. And then Vp is delta to the T times Vm. So if you look at what's going on here, these are just-- you can think about bringing all of the Vm's and Vp's over to one side of the equation. So, for instance, this first one looks like Vm minus 1 minus alpha delta Vm.

And then Vp is delta to the T Vm. So this is really minus alpha delta to the T plus 1 Vm equals 1 minus alpha times pi m over 2. So just by substituting in for this one, I just get an equation in one unknown. I can solve that for what Vm is. But it's just a two equation and two unknown problem. Anyway. So this is what I get. So this is the expression for what Vm is, and this is an expression for what Vp is.

So then what we're trying to show is that this is true. OK? Again, I bring this term over to that side. So what I need to show is that Vm minus delta to the T plus 1 Vm is at least pi m times 1 minus alpha. So that's what I wrote here. The Vm has a 1 minus alpha in it. I canceled the 1 minus alphas. This term also has a pi m in it, as does this term, so I cancel the pi m's.

I multiply out. Anyway, I get-- a bunch of algebra. I get this expression. 2, 1 minus alpha, delta plus 2 alpha minus 1, delta t plus 1 is greater than 1. So if you look at this expression, if alpha is less than a half, 2 times 1 minus alpha is bigger than 1. So I have a number bigger than 1 times delta plus something with a delta to the T plus 1 has to be bigger than 1. So if 2, 1 minus alpha is bigger than 1 and I choose delta close to 1, this is bigger than 1.

And then I can choose capital T large enough to make this thing approximately 0. So I've got something bigger than 1, plus something that I can make as close to 0 as I want. That's going to be-- that is going to be greater than 1. So just the intuition is just make delta big. Make delta large so people care about the punishments about the future, and then make T large, which is just making the punishments as severe as you need it to be. So if delta is close to 1 and T is large, this is going to work.

And I described that as T going to infinity. But if you think about the math of it, if I gain, I gain in one period, I gain half the monopoly profits. And then in every future period, I'm losing half the monopoly profits. So even losing half of the monopoly profits twice may be enough to make me not want to get half the monopoly profits once. So capital T equals 2 will often work for this. Capital T equals 3 will almost always work.

There is problem with the 1 minus alphas in there, that you may not actually get punished here, because it may be that there is no demand. But if alpha is small so that the punishment really does occur, you just don't need T very big to get this thing to work. OK, so while it's a very artificial model, I think it's a very insightful one because it makes several important predictions. So first one is maybe unintuitive.

Price wars should be observed in equilibrium. I think there was-- if you read the old literature, there's this very nice book by Tom Ulen about railroad cartels written in the 1970s that's thinking about these cartels tried to organize collusion, and the collusion kept breaking down. And they kept having price wars because they couldn't collude with each other. Green and Porter makes a point that, no, price wars are not something that's a sign that the cartel has broken down.

If you have a well-designed cartel, it has to have these price wars in it. Because if it didn't have the price wars in it, people would be cheating. And so when we see price wars, we can think well-functioning cartel rather than poorly functioning cartel that keeps breaking down. The optimal collusion in this model can have T finite. Why is that? Certainly, making T infinite makes it the biggest punishment and makes people not want to deviate. But if you make T infinite and alpha's not super small, then just basically only collude at the very beginning of the game.

And so what you want to do is have a punishment that keeps people from deviating but not so long that you lose a lot of revenue while you're in the punishment. So the best equilibrium at this point is going be the one that has the shortest T possible that deters collusion. So we should have these short, finite price wars. What causes price wars? In this model, the insight is price wars are caused by demand patterns that make the participants think someone has cheated.

In equilibrium, no one ever cheats. It's not people cheat, price wars start. It's always suspicious demand patterns arise, people start price wars knowing that it was-- with probability 1 and equilibrium, it was a suspicious demand pattern. But it's suspicious demand patterns that trigger price wars among firms. And you might think, well, if the firms all know that it's a suspicious demand pattern that triggered it, can't they just renegotiate out of it and say, OK, look, let's go back and ignore that one because we all know the firm made no orders this week?

Once you do that, it doesn't work. The firms have to commit that we're going to have these price wars when these suspicious demand patterns arise because it's the suspicious demand patterns that resemble someone cheating. Yes?

AUDIENCE: Are there cases we know about where we know it was-- the price war was due to the cartel functioning and not just demand?

GLENN ELLISON: Yeah. So I will-- in Wednesday's lecture, I'm going to talk about Rob Porter has this paper re-examining a railroad cartel from the 19th century. And there-- I don't know. I guess he at least-- he argues very much that this model looks like it's describing what's going on. I think it is an open question, in that example, whether it's cheating or whether it's suspicious demand shocks or whatever. There are plenty of cartels where we know that someone cheated and that someone cheated and that triggered price wars.

The no one cheated is often harder to know empirically because the whole point is that these things are unobservable. And so how do you prove that no one did cheat and caused that shock that made people, everyone start the price wars? But actually, it's interesting. I will talk about this next time. But in the-- collusion wasn't illegal in the United States until 1887. And so in the 1880s, there were collusion consultants.

So just like firms can go out and hire McKinsey today, in the 1880s, you can go and hire a collusion consultant and say, we have our industry. We're competing too much. We're not making-- this was the era of the Vanderbilts and whatever. We're not making enough money. My mansion in Newport only has 137 rooms. I need higher prices. And you would hire-- there were McKinsey-like firms that you would hire, and those firms would consult with you on how to set up your collusive agreement so that you would be able to keep prices up or whatever.

This model is very stark, that you get these random-- no one ever gets away with cheating in this model. In a model with continuous demand shocks, you would get away with cheating. So let me say Green and Porter do a different version-- do a better version of this model. Green and Porter, they do a Cournot version of this model, where the firms choose Q1t and Q2t. And then you have price at time t equals P of Q1t plus Q2t plus epsilon i t. So you have something that looks like real demand shocks.

The firms choose quantities. There's some, then, downward-sloping demand curve. But then demand is also just random in every period. The market-clearing price is random. If you think about this, there's some price they're trying to coordinate on the monopoly price as the outcome. But you're going to get this distribution of prices even in equilibrium. And so what would you do if, every period, you're going to get a price drawn from this normal distribution around Pm if everyone's behaving properly?

You can't say, OK, if price isn't Pm, we start a price war because that would happen all the time. So what you do is you do some cutoff here, like p-hat. And if price is greater than p-hat, we keep colluding. If price is less than p-hat, we think it's likely enough that someone cheated that we start a price war. In a model like this, I can cheat and not get caught.

If I cheat, I'm pushing the price distribution down. But pushing the price distribution down, there's still this large probability that we don't start a price war. But I do increase the probability of the price war. And you have to hope that the-- it's got to-- so the increased probability of a price war from a small deviation is enough to get you to not make that small deviation. In a more realistic model, firms will-- firms could cheat and get away with cheating and only get caught probabilistically.

This is part of what also makes it harder for us to know, was there really cheating there? Because if you do epsilon cheating, there's only-- if you have an order epsilon cheat, you get an order epsilon increase in profit and an order epsilon increase in causing the price war. Hard to know if that's happening or not. Yeah?

AUDIENCE: Sometimes, in this model, you will get away with cheating. Do you mean in the sense that there's a mixed strategy in which you cheat sometimes or in that the equilibrium--

GLENN ELLISON: No. So the equilibrium we're going to look for is an equilibrium with no cheating. It's an equilibrium where firms always charge exactly these prices. But if you did go to Q1t minus epsilon, then you'd just raise the probability of a price war by epsilon. So there's some probability, x star, that we should start price wars. And now price wars start with probability x star plus epsilon or something. Yeah.

And sustaining collusion is more difficult when demands are noisier. Here in this model, when alpha gets large, you can't-- you need alpha less than a half to make this work. In this model, if the price distribution looks like this, it's relatively easy to collude because you just put a threshold here. And no one wants to shift demand down below that. But if prices look like this, it could be that, no matter where you draw the cut-off, the increase in probability of a price war from cheating is so small that you're not pushing much mass beyond any one of these points.

It's hard to stop collusion. It's hard to collude. So the closer we get to that first model example of you see someone cheat, you know they cheated, the easier it is to sustain collusion-- and the more there's noise in that, the harder it is. Questions? So obviously, this was a very beautiful 1984 paper that first did this. As I said, this is how Green and Porter actually did the model. It's harder to solve but more economically sensible.

There's a big subsequent game theory literature. If you take [14.]126, they cover many of these models. One very nice paper on this-- Abreu, Pearce, Stachetti, *Econometrica,* 1990-- discusses more general strongly symmetric models like this. What Abreu, Pearce, and Stachetti show is that in the more general version, like Green and Porter's, you actually want to have-- they give this dynamic programming thing where it is basically always a two-- always a two-phase thing.

But you actually-- what you end up doing is you do punishment. You do incentives on both sides. Here, firms don't want to cheat from the high price because if they cheat, we start a price war. But on the price war phase, what gets firms to charge these really low prices is the lower is the resulting price. You only switch back-- you switch back to the top not when-- not after T periods, but when price ends up low enough. And so it's like the firms are trying to drive price low enough to get back to a price war-- get back to colluding again.

And so, in some sense, you use the incentives on both sides. Here, they have an incentive not to cheat because it causes this. Here, they have an incentive to price low because the lower they price, the more likely it is to jump back up. And so you still get this two-phase thing going on, but you get this two-phase thing where you have incentives. Rare events cause us to jump down, and then rare events cause us to jump back up, and that's the optimal collusive agreement. Yeah?

AUDIENCE: Is the reason why they're pushing the price even below cost just to make the punishment painful enough that everyone goes back up? Or why would that happen?

GLENN ELLISON: Yeah. Well, if you think about it, punishments are wasteful, right? You're in a punishment phase. You had a period where there was demand, and you served at a profit zero. So wouldn't it be better to lose a little bit more money today? Because your D pi D P evaluated at c is small or whatever. So you lose more money today, and then you can get more money back in the future or something like that.

So, just in general, you want to give firms as much incentive as you can to make them low. And a good incentive to price low is that things will get better in the future if you-- the lower your price today, the more likely it is things will get better in the future is an extra incentive to get them to carry out the pricing low. Yeah. I'm not sure if we need it in that p equals c example with pure Bertrand competition.

But in general, you want the incentives on both sides. i guess a thing about-- Abreu, Pearce, Stachetti also writes that, in some sense-- what they'll show you is that it's wrong-- if you're trying to think about what are all the possible strategies you can use, it's generally a bad idea to focus on strategy space because strategy space is very, very complicated. What you want to do is focus on payoff space, where these things are just numbers, and think about what's the highest Vm we can use. And to support the highest Vm possible, we want to use the lowest Vp possible.

And to just think about what constraints the Vm and Vp have to solve and think about that from a dynamic programming perspective, just what's the highest Vm, what's the lowest Vp. What constraint would Vp have to satisfy? What constraint would Vm have to satisfy from not wanting to deviate? And then use that to-- I don't know. I would say that's a big insight from Abreu, Pearce, Stachetti, is we want to focus on this payoff space rather than on strategy space. Because payoff space can be just two dimensional, whereas strategy space is infinite dimensional.

OK, then. Fudenberg, Levine, Maskin, a few years later, have this paper showing that, in many models, we can avoid the inefficiency that's in Green-Porter and Abreu, Pearce, Stachetti by using asymmetric strategies. The idea here is that let's suppose that you actually get a signal not only of that someone cheated, but who cheated. Now, the Green-Porter model, you can't tell who cheated. If this is firms are supposed to be choosing these quantities, price is too low, there's no common signal of who cheated. There's just that someone cheated.

If the demand was at all asymmetric with regard to Firm One's price and Firm Two's price or Firm One's quantity and Firm Two's quantity, then you would get a signal about who cheated. And so what Fudenberg, Levine, Maskin say is that, well, if you actually had a signal about who's cheated, you should get rid of the price wars, and you should just do, say, a market share adjustment. If Firm One prices-- if it looks like the demand makes me think Firm One cheated, then I should have, instead of punishing Firm-- instead of having a price war, we should just have a market share reallocation, where, in the next period, Firm One gets a smaller market share and Firm Two gets a larger market share.

And by doing market share reallocations, we avoid all the inefficiency of the price war, and we still punish people. So, for instance, if you assumed there were these two firms, one of them had larger demand on the Eastern US and one had larger demand on the Western US, then if you saw prices go down more in the East than in the West, you'd be like, OK, the East Coast firm cheated. Therefore, next period, East Coast firm has to charge monopoly price plus something. And the West Coast firms charge monopoly price minus something, and we reallocate market shares so that the East Coast firm is relatively disadvantaged but where there's no first-order loss in profits.

So in general, that's the Fudenberg, Levine, Maskin ideas. If there is an asymmetry-- Green-Porter is unusual in not having any asymmetry. But if there is an asymmetry, you just want to reallocate market shares between the firms rather than doing that. Anyway. So there's one place here where there's recent literature, is the private monitoring, where-- all these papers are focusing on this public monitoring where everyone observes the same signal. Everyone observes the same signal of what happened.

If you imagine-- if you had a model where the firm said P1t-- so we have these firms set. P1t. And then you get demands Q1t equals some function Q of P1t, P2t plus epsilon one t and Q2t, Q2 of P1t, P2t plus epsilon 2t. This would be what we would call a private monitoring model, where each firm gets a different signal. So Firm One observes its demand. Firm Two observes its demands. And they have different signals because they have different epsilons.

This would be a private monitoring model where I don't know what signal you got and you don't know what signal I got. Anyway, it's harder to solve models once you have that private uncertainty instead of the public uncertainty. There's a literature on that. Actually, Alex has a very nice, recent paper arguing that sometimes, actually, private monitoring will make it more complicated. Could even make collusion easier.

Although, obviously, sometimes, the normal case would be it makes collusion harder. But some cases where that can even make collusion easier. Any questions? OK, the next set of things I wanted to just-- I'm going to go through and highlight them quickly. Maybe I'll come back to them at the end. But I wanted to get to more recent material at the end. So just some classics. Rotemberg and Saloner, 1986, argues that variable demand can make collusion harder.

Let me just say, the basic argument here is-- as I said, it's harder to collude when the present is more important than the future. And when the future is more important, it's easy to collude. When the future is less important, it's hard to collude. And so the basic intuition for Rotemberg and Saloner was that imagine we just have firms competing, and they're just facing a macro economy that has ups and downs to it. So if you just have business cycles like this, where this is t and this is Qt of p star-- so imagine you have firms that are trying to collude with their business cycles.

And so there are times when they're selling-- 1955, they're selling a lot of cars. 1957, they're not selling as many cars. You know, Rotemberg and Saloner note that, in an economy like this, it's harder to collude when demand is variable because here the present matters a lot relative to the future, whereas here the present future would matter more. So here, it's relatively easy to collude. Here, it's hard to collude. And so, then, what Rotemberg and Saloner ask is, well, what would firms do if they were colluding in a market like this?

And imagine the delta is such that you can't collude on the monopoly price here, but you can collude on the monopoly price here. So what do you do in these periods? So if there was some-- imagine this is just a multiplicative shock to demand. So what I'd like to do is have the same monopoly price in every period. I'd like to have the same monopoly price in every period, but what I'm going to need to do is, when we're in these value of the high demand shocks, I'm going to have to drag down my price and price it less than the monopoly price.

And I'm going to have to price it less than monopoly price because I just pulled my price down to reduce the incentive to-- temptation to deviate. And so what Rotemberg and Saloner argue is that, if you have a cartel, it's harder to collude during the booms. That doesn't mean you should have price wars. That inclusion should break down. But what you should have is this countercyclical pricing, where, because demand is high, we drop our price down. And we could actually see lower prices here than we see here.

Now, obviously, depending on what demand specification you make, if I do linear demands-- in linear demands, price goes up when demand goes up. So it could be that the monopoly price isn't a flat line. It could be the monopoly price would also be something that mirrors the demand. Then what Rotemberg and Saloner say is, OK, you're going to pull down your price here. But pulling down your price here, it could still be the-- it could still be higher than the price here because it's-- I've pulled down my price, but I pulled it down from something that was starting even higher.

But anyway, the basic idea of Rotemberg and Saloner is that you can get this countercyclical pricing that comes from a countercyclical incentive to deviate on a cartel agreement. Anyway, on the slides, I've worked out an example in a two-state model where demand is just i.i.d. high or low. And when demand is i.i.d. high or low, the algebra is easy, and you can work out what you do. But you can look at that on the slides, if you like.

Let's see. Which of these things do I want to say? Yeah. So for intermediate, markups are countercyclical. Yeah, If you add imperfect observation to this, intuition is still that collusion is harder when demand is high than when demand is low. And so what are you going to have to do? One of two things has to happen. Either you have to reduce the prices and reduce the markups, or breakdowns of collusion become more likely, or both.

But in general, you could adjust either of two ways. You could either drop the prices, or you could keep the prices up and just have more frequent breakdowns of collusion. Or you could do a combination of the two of them. Or you could do one, but not the other, possibly, depending on the specification. But it's just the general idea should be that collusion is harder and something has to work worse.

Second idea I'll do quickly. Private cost shocks. So Athey and Bagwell note that efficient collusion you can think of as having two components to it. One is that firms want to charge the monopoly price. The other is that the colluding firms want to produce efficiently. And in some cases, the firms are going to have cost shocks that make it efficient for one of the two of them to serve.

So, for instance, for this one, for an example, think about-- let's imagine-- and this may not be so unrealistic-- that in a city like Cambridge, if there's a giant paving job that needs to be done-- you need to repave Memorial Drive or whatever-- there's a very small number of firms that could do that job. And so what the city of Cambridge is going to do is put out for bids. Like, we need someone to repave two miles of Memorial Drive. There's a very limited number of firms who could do that job.

Those two firms are going to have different costs of doing the job in each period because their equipment may be in use elsewhere. It may be that I've got a job in Revere where I'm paving some big road, and it wouldn't be costly for me to also get extra manpower and equipment to get to Cambridge. Or it could be that I don't have any other jobs, and so this would be a good time for me to serve Cambridge. So what they assume is that-- we have these two firms. They're serving some demand.

And the firms have random cost shocks. Again, to make it simple, they just go with two different levels of cost, C L or C H. And what the firms would like to do is collude on charging monopoly price and collude on whichever of us has low cost wins the job and whoever has high cost doesn't win the job. So in some sense, what we should do every period is-- so the city of Cambridge just repeatedly issuing these requests for bids.

If our costs are C L, C H, what we should do is I bid the monopoly price and you bid the monopoly price plus epsilon. And if the costs are C H, C L, then we can bid monopoly price plus epsilon monopoly price. And if the costs are C H, C H, we can just go ahead and bid monopoly price, monopoly price. And just one of us is chosen at random to win that auction. So this is what we would like to do.

How do you do that? Well, the firms have this other problem now. The problem is this communication between them. The question is, who gets to bid the monopoly price, and who's supposed to bid the higher price and let the other firm thing? In some sense, every period, the firms can get together and be like, OK, who wins this job? Who has the C L cost? And everyone's going to raise their hand. Oh, yeah.

My cost is C L this period. I should get this job. Bad luck for you, but I guess you bid price plus epsilon or whatever. And then, again, you have this temptation, over and over and over again, the firm would raise their hands and be like, oh, yep, me, just happened to have all my equipment, a bunch of spare labor. I need to win this job again this month. What Athey-Bagwell shows, that's an additional obstacle to collusion, this desire to-- that having the need to shift the market shares across firms is an extra problem.

One thing they show is, if delta is close to one, the firms can still achieve full collusion. And this isn't a problem. The way they can do this is with this dynamic allocation of market share. So the rule could be, if my costs are really low and your costs are high, I go ahead and bid Pm and I win. But then the next time C L, C L happens, if you said your cost was low and you wanted to be the winner of this auction, then the next time an auction comes around where we both raise our hands, you have to set the high price, and the other firm gets to set the low price.

So it's like you dynamically make up for the market share gains that firms got from claiming their costs was low by saying, OK, if you claim your cost is low today, then the next time we both claim our costs are low, I get the demand, and you don't get the demand. And so the idea is that, by dynamically reallocating market share, we can sometimes collude optimally. When delta is not as close to one, however, this stops working.

And the problem is, you can get this same problem I discussed before of just some firm, over and over again, says, yeah, my costs are low, my costs are low. My costs are low. You can't keep promising, OK, next time your costs-- next time both costs are low-- in some sense, you're accumulating debts to the other firm. And you can't accumulate too many debts before this thing stops working. So at some point, you would have to get to an agreement where, OK, if you've done too many times in a row, even if your costs are low and my costs are high, you still have to back out and do this.

Or we can stop you from deviating-- stop you from doing that by saying you've done it too many times in a row. If you claim your costs are low, you have to cut your price below the monopoly price instead of charging the monopoly price. Because we need some extra disincentive. Once you're already saying, in the year 2045, I'll give you market share, that doesn't work. So you have to-- for the firm, if you're going to say your costs are low, we're going to make you cut your price.

So anyway. I think the message of Athey and Bagwell is that it's the cost shocks are a problem, but they're-- if firms are clever and can do these dynamic reallocations, because it's frequent things, that may not be such an obstacle after all. Harrington has this paper about antitrust authorities. Obviously, in practice, we do think that one of the reasons firms don't collude is because it's illegal and you can go to jail. And business executives who make millions of dollars don't like spending time in federal prison.

It's unpleasant. Obviously, so that is a constraint on antitrust, on collusion is that people can go to jail. And firms can pay very large fines. One thing that Harrington notes is that that is true. But you can imagine reasons that doesn't work so well. So, for instance, once you have those penalties there, firms really don't want the penalties to get triggered. And it may be that something that triggers a penalty is a breakdown of the cartel agreement.

So if we have a price war and the price drops from 100 to 60 and then goes back up again, that's a clear sign to antitrust authorities that maybe we should look into what's going on here and are they colluding with each other. Therefore, you may not want to cut price and cheat from the agreement in the first place. And therefore, you may-- in some sense, you can drift into collusion and get even more collusion with the antitrust authority than you can get without because no one wants to risk starting the price war that then makes people notice that there was a cartel agreement going on.

I think while I do think the main version of that is certainly true, that collusion being illegal is a deterrent, it can also be-- not well-designed collusion policies can also, then, in some sense, lock people into cartels. Something you often see in practice reflecting this is that, sometimes, these things have whistleblower protections, that if you go to the government and say, hey, I've been colluding with these guys for the past five years, you get off with a much lighter punishment than everyone that you were colluding with. And that can be a way to try to get rid of this problem.

Multimarket contact. Limited time. I will leave that for you to read, also. So things I wanted to do. How do people in IO get away from this indeterminacy and not just always have this conclusion that we can't do empirical work or we can't do other things because we can't tell what's going on? One thing we do is we adopt this static Nash framework of just, OK, we're just going to assume static Nash described what goes on.

In a lot of models, you can't just assume static Nash is what goes on because there is some variable that's changing over time. Like, firms are investing in their equality or in reducing costs. Or firms are setting prices, and it's costly to change their prices, so whatever price they set today are locked in for the future. We need some model where firms-- there needs to be some trade-off between today and the future to think about staggered price-setting or to think about cost, quality investments, or whatever.

What people often do is to use this concept called Markov perfect equilibrium. In a Markov perfect equilibrium, players' strategies are assumed to depend only on payoff-relevant variables. So, for instance, you can price on-- my price at time t is a function of my cost and your costs. These things are relevant to the payoffs in the game. And therefore, if these things are changing the payoffs in the game, I'm allowed to price based on this.

But I can't be doing I'm pricing today based on the price I set at t minus 1 and the price you set at t minus 1 because those prices are completely irrelevant today. So in Markov perfect equilibrium, what we want to do is assume that firms can set their strategies based on payoff-relevant variables but assume that they can't set their price based on payoff-irrelevant variables. So saying that can't happen. This is what rules out those grim-trigger punishments in the repeated Bertrand game.

You can't say, if you said price Pm minus epsilon yesterday instead of Pm yesterday, then we go to a price war forever. I'm basically trying to recreate-- get as close as I can to static Nash by saying, you can only price based on things that are payoff relevant, not based on things that are payoff irrelevant. And people often think of this as a way to get back to obtaining unique equilibrium. We see an interesting observation is that people often think of Markov perfect equilibrium as doing that.

But in practice, it doesn't always do that. And so there's this-- I talk about this interesting, beautiful paper by Maskin and Tirole that just give examples of what kind of Markov perfect equilibrium could we have in an alternating-move price game. So in Maskin and Tirole's game, it's just this-- it's this alternating-move game. We have two players. Let's say Player 1 and Player 2. And Player 1 picks some price at, say, t equals 1 and sets some price that's going to have to hold for two periods.

And then Player 2 gets to pick a price in period t equals 2, and that has to be there for two periods. And then Player 1 chooses a new price at three that holds for 3 and 4. And Player 2 chooses a price that holds for 4 and 5. But the idea that we're trying to get at is, in reality, price changing is asynchronous. If you look at two gas stations across the street from each other, it's not like the guys climb on the signs and change the price listings at exactly the same time and change the pumps at exactly the same time.

It's like somebody changes their price, and then the other guy sees the first one's price and changes his price. And they will tend to go back and forth like that. Each one observe it-- needing some time to observe one's price, do something [? rather, ?] and then change their price. So they consider this stylized game where Firm One chooses prices at t equals 1, t equals 3, t equals 5. And Firm Two chooses prices at t equals 2, t equals 4, t equals 6.

And then, besides that, they do repeated Bertrand competition. So just repeated Bertrand competition. D of p is 1 minus p. And to make a simple example, they restrict prices to this simple grid-- 0, 1/6, 2/6, 3/6, 4/6, 5/6. One observation is that, in this model, there's no Markov equilibrium in which the firms price at cost. So the Bertrand equilibrium of just pricing at c goes away. And why is that?

Well, because imagine you had an equilibrium where the firms were supposed to price at cost in all periods. What Firm One could do is just, in the first period, instead of pricing at c, I raise my price and I charge. So suppose c was 0. Instead of price 0, I raise my price to 4/6. I set a really high price. Then in period two, is Firm Two going to continue pricing at cost? Well, if it continues pricing at 0, it earns zero profits.

You've given it this opportunity. It knows you're locked in at this really high price. It could raise its price to 3/6 and charge them at-- which is the monopoly price in this model and earn the monopoly price for two periods. If it does that, then, here, I can price at 2/6 and earn some money, and I'm earning 0 in equilibrium. So we can't have a zero-profit equilibrium here because firms could always raise their prices, tempt the other firm to raise their prices, then undercut that other firm and earn something positive.

So zero-profit equilibria go away in this model. They don't exist. What does exist? One thing that still exists here is collusive agreement, collusive equilibrium, like the ones that Markov equilibrium is, in some sense, designed to get rid of. So consider this collusive equilibrium where, basically, we always, whatever the other firm is charging, if it's high, we charge the monopoly price of a half. And then if someone undercuts the other firm and goes to 2/6, then you undercut them and go to 1/6.

And then if the other firm is charging 1/6, you either understick at 1/6 or you go back to the monopoly price. So this is like a-- you can think of this as a probabilistic price war. We're going along. We're supposed to be charging this-- be Pt-- we're supposed to be charging this price of one half in every period. If someone deviates and instead goes here, then the next firm, the firm goes there. And then they have this randomization between if they do that. And then somebody gives up and goes back to one half, and then we continue.

So what they basically do in this Markov framework is, because the other firm's previous period price is payoff relevant, you have to include it in the Markov state space that you price on. And so we can get these price, probabilistic price where we're looking just like Green and Porter, then jump back up. But this is also a paper-- like a folk theorem paper, there's also equilibrium multiplicity here. And we get these other equilibria too.

So one interesting one they point out are these things called Edgeworth cycles. Edgeworth cycles, we never settle down on pricing. And what the firms just do is forever undercut each other. So the firms just do this. They start out at a high price of 4/6, and then they just undercut each other. They get down to the bottom. There's some mixing. And then somebody wants to end. Again, this is like incentive to end-- you have this incentive to end the punishment, where it's like each one wants the other guy to end the punishment.

But then someone jumps up to 5/6 and lets the other one undercut him in the next period to 4/6. And then we just continue like this. So we can get these Edgeworth-cycle equilibria where the firms just continually undercut each other over and over and over again. And then, at the bottom of the cycle, you've driven the price all the way to 0. What do you do? At some point, you give up and say, OK, we've been at 0 for long enough.

I'm going to go ahead and price at 5/6, also earning nothing, but then let you undercut me to get to 4/6. And then I'll get to 3/6 two periods from now. So we also have these Edgeworth-cycle equilibria. And so I think the lesson here with this paper was that we may think of Markov equilibrium as eliminating multiplicity, but if-- here, it's because the previous period's price is payoff relevant, and that would happen in any staggered price-setting game. Or you could also do this based on other things that are payoff relevant.

Like, if costs are there, we can also bootstrap on the costs. But so what we can get is multiplicity, and we can actually get multiple interesting patterns. We can get the collusive pattern we've seen before. You could also get this kind of pattern as an equilibrium. And the one thing you can't get in that model, interestingly enough, is competitive pricing. So competitive doesn't work. But this is possible.

Sometimes, it's funny that perpetual cycling of prices makes more economic sense than just static Nash pricing in this model. Any questions on that? Obviously, this is a simple numerical example. But if things happen in this numerical example, similar things are going to happen in lots of more complicated models, that you can get these multiple types of Markov equilibria. So then the final thing I was going to talk about today is a more recent paper, Calvano, Calzolari, Denicolo, Pastorello.

This is a 2020 *AER* paper motivated by thinking about online businesses. And they note that, in many online businesses, firms are setting prices repeatedly. And we really are in a case where delta is very, very close to 1. If you think about Hertz competing with Avis, competing with Budget on selling rental cars through Orbitz, they can change their price every second. They have a computer monitoring the other prices and changing prices. You think about competing sellers selling through Amazon.

Many firms are trying to sell umbrellas through Amazon. They're competing with all the other firms trying to sell exactly the same umbrella or the same Nike shoes through Amazon. You can also think about Amazon competing with Target. There, the prices are not appearing side by side. But it's still Amazon choosing the price at which it's going to sell something, Target and Walmart choosing the price at which they're going to sell it, and the demand is a function of all those prices.

And again, Amazon is not having a guy decide what's the price at which I'm going to sell this pen. Amazon has got a computer that's monitoring the internet and changing prices continuously and trying to figure out what a good price is. And so what the paper's asking is, what do we think is going to happen in that situation where you have computers setting prices, competing against other computers in a repeated-game setting? We know that there should be many, many equilibria. In some sense, the question is, which equilibrium will the computers find of the many repeated-game equilibria that exist?

And they notice that we should-- when we think about what the computers are going to be doing, the computers are going to be trying to learn what's happening. The computers are not born knowing, what's my demand when my price is 39 and your price is 38? The computers have to figure that out. And the computers also have to figure out what's your strategy. If you are at 39 and I undercut to 38, do you undercut to 37? Or do you stick with 38?

The computers, in some sense, have to figure out what should we be doing to profit maximize. And so they try to think about what would happen when you have agents who are trying to learn competing with each other. And in some sense, this is resurrecting-- there was a large theory literature on this in the 1990s. Drew has this very nice book with David Levine from 1998 discussing lots of theoretical results about what happens if people are naively trying to learn.

And at the time, the theorists were arguing that this is a very interesting question because if people are naively trying to learn, do they learn an equilibrium or not? To believe in equilibrium, we need to think that that process converges. And if that process does converge, then the question is, well, if we have a game with multiple equilibria, are there some equilibria that are more likely to emerge from that learning than others? So, in some sense, that's what Calvano et al. are trying to do here.

So they bring in a more complicated learning model than existed in the 1990s theory literature. But it's a learning literature, thinking of how operations, research, or computer scientists would say you should design a pricing algorithm for your website. So we have a situation where profits pi depend on s-- a state of the world that's observable to the firms when they're setting prices-- and a-- an action that they take which you can think of as their price. And the profits depend on this observable state, this action, but it also depends on a random shock.

So I set my rental car price at $39. Someone else sets it at $38. And then look at demand over the next day, there's some random number of people who reserve a car at LAX on that date given those prices. And then there's some state transition process, probability that the state goes to s prime given s and a. That's also unknown. So again, the staggered price-setting game, you can think of the s, the state of the world, as everybody's price on that rental car website.

The current prices today out of LAX are $41, $40, $40. Something like that. That's today's state. I observe the prices are $41, $40, $40, where I'm the third of those prices. I choose what price do I charge. And then I'm going to get some profit over the course of the next day. And then, the next day, when I go back to check the website again, all the prices are going to be different. It may be $35, $38, $39.

And then I'm going to try-- sometimes, I don't know what that state transition process is because I don't know what my rivals are doing. And I also don't know what my profits are because I don't know the demand curve. So what would you do? Well, dynamic programming suggests that a good thing to do is to try to learn this function, Q of s, a. Often, we talk about this value function, the value that you get being in state.

I'm going to find v of s, a. It's going to be the max over a, Q of s, a. Sorry, that's behind the board. So Q of s, comma, a is the profit you get if you're in state s and you play a today and then you play optimally in all future periods. So Q of s, a, the recursive definition of it is Q of s, a is the expected profit you get from taking action a when you observe state s plus delta times the expectation over what tomorrow's state is going to be given that state was s and you took a of the maximum over all actions you're going to take next period, Q of s prime, a prime.

So I'm going to take action a. And then, next period, I'm going to see some new state s prime. And therefore, I'm going to choose some new action a prime. And then I'm going to get this discounted present value of profits, Q of s prime, a prime. So if I knew this object, Q of s, a, then I can behave optimally, right? I just do the argmax over all-- I observe a state.

I just choose the a that maximizes Q of s, a. So if you think about it, Q is just a matrix of numbers that I need to know. I have all the possible states-- s, s prime, s double prime, s triple prime. And there are all the possible actions I can take-- a, a prime, a double prime, a triple prime. I just need to know what's my discounted present value in each of these-- in all of these things. If I know the numbers in this matrix, I just-- whatever state I'm in, I just maximize the numbers in that column, and that's how I behave.

So here's a standard learning algorithm for how you might program a computer to behave in this situation. You conjecture some initial matrix Q. So maybe you have some initial conjecture demand curve, and maybe you assume just some flat conjecture about how your rivals are going to price and figure out, what profit do I think I'm going to get if there's no dynamics? I just guess what the demand curve is.

I just initialize Q to some-- I initialize Q to something. So I assume that this is the profit-maximizing thing. So my profit is 2, 2, 2, 2, 2. 1.9, 1.9, 1.9, 1.9. 1.8. I just initialize this matrix to something, guessing that I think this is the best and I think my profits stayed independent. Then, at every point t, I figure out, OK, if this is the state at time 0, that's my optimal action. So I take it. So I choose the profit-maximizing action given my current knowledge of the Q matrix and my current observation of the state.

Well, I choose that with large probability. But with epsilon probability, I experiment because I don't know. It may be that this price is not good. So I usually choose this. But with some probability, I experiment and choose a different action instead. And then when I choose this action, I update the number in this cell by saying I'm going to take a weighted average of what I thought the number was-- so I thought the number was 2.

So I take a weighted average of 2. And what I now-- what my current period experience is, which is I got profit pi t, and then I got profit pi t and then I end up in this state, and I think, what's the best I'm going to do in this state? And so I do-- I know that I got profit pi t, and I know that I ended up in state s t plus 1. So I do this plus what I expect from optimal play in that state, still using my Q matrix that I started with.

So here, actually, I think I'm still going to get 2, even though the state has changed. Anyway. So theorem from people who study these kinds of learning algorithms is that, if I'm playing against a stationary environment, there really is just some exogenous process by which states transition, which can depend on my action. And it depends on the current state. But if there's some Markov state transitions like this and there's some profit function-- I don't even know what it is-- this Q-learning algorithm finds the best thing to do in the long run.

Again, you need some of ability of full support of the Markov process so you don't get stuck. Like, there was some great thing to do, but once you leave that, you can never get back to that state. But if there's-- if the Markov process is going to cover all states, this thing will eventually find it. Because you're just updating-- every period, you're just updating one element in this to reflect the profit you get in this state and then the profit you get in the next state.

And then if you're in this state, you update the profit here to reflect where you end up there. You just keep updating the profits, taking in a weighted average of what you thought they were and the current thing. So at some point, the current, the actual experience starts to overwhelm the past experience. This thing converges. And so Calvano et al. conjectured that this would be a reasonable thing for firms to do if they're pricing against another firm on a travel website.

If I were to assign you-- like, here, I want you to design a pricing algorithm for Hertz to use, that you might have come up with this. OK, so what they do is then they study what would happen here in a repeated competition game. So they go to things that look very much like we're used to. So they go to a logit price competition game where consumers have these v minus p plus epsilon preferences. And the epsilons are logit distributed errors.

Their base model has two firms in it. The two firms have identical costs. They have a grid that has 15 possible prices. They have a discount factor of 0.95. They don't know what parameters firms would use. Alpha, if you remember, was the weight on current profits, how fast you update the numbers in the Q matrix when you get profits in each period. And again, here, profits are random because the customer-- because the epsilons might buy from you, might not buy from you.

And again, you don't know the logit parameter, but you just know that higher price less likely, but I don't know the elasticity. And then the second parameter beta, which governs how fast the experimentation rate goes to 0. As I said, you're going to want the-- early in the game, you're going to want to experiment a lot. After you've been playing this thing for a million periods, you want to stop experimenting and just start picking the best thing. And so these two parameters, alpha is how fast you adjust your Q matrix, and beta is how quickly you abandon-- how long you abandon-- how slowly you abandon experimenting.

Anyway. They declare learning to have converged when optimal play is unchanged in their model for 100,000 periods in a row. Obviously, they had a big computer. And they also, actually, consider 10,000 different parameter combinations, alpha and beta, to be able to talk about what happens here based on what alpha and beta do firms use. Main results they have are that, really, it turns out that, for a wide range of alpha and beta, pretty much the same things happen. What happens?

First, profits are usually 70% to 90% of the way from static Nash to monopoly. So what we find the computers learning is to collude with each other. And so they end up with these collusive pricing that's pretty far towards pure monopoly from static Nash. Where do the computers end up? The computers end up playing things that are epsilon Nash equilibrium. In some sense, the computers invent punishments with finite price wars.

So the computers develop these beliefs about what's going to happen. That means that we all charge this monopoly price. And then if-- so let's suppose this is the state where we're both charging price two. And then what the computers think is that, if they deviate to 1.92 and we're in this state s prime, then in the state 1.92, profits are lower. And profits are lower in a way that's going to make the state-- the other firm undercut. And we're going to transition and follow this arrow in the matrix.

In some sense, the computers learn that, from here, profits are going to be such that I want to undercut and charge an even lower price. And then, the next time, he's going to want to undercut and charge an even lower price. And the computers, in some sense, learn that the payoffs are such that you want to do something not so different from one of these mass [? interal ?] cycles off the equilibrium path.

So even when you set delta approximately equal to 0, the computers still learn that they shouldn't-- delta is positive but nonzero. The computers still learn they should collude. They don't collude much when delta is small. They collude with 20% of the weight of the monopoly price. But it seems that, whatever you pick, the computers are colluding that way. And then, as you increase the number of firms, it does decrease the collusive level of prices.

So if we want to ask what causes prices to go down as delta increases, like I said, the pure theory version has a very hard time with that because 1 minus 1 over a quarter is not a big number. But this says that prices do seem to decline the number of firms, and they learn less collusive equilibria when the number of firms is larger. But still, high with four firms. And let me say, it's hard to write papers like this because you're used to-- in a theory paper, people want to see theorems.

And you can read this as an example for, how do I present results from a paper that has lots of simulations instead of theories? And they have a lot of nice, pretty graphs in there, and I think that helps. So this is their profit gain as a function of the two parameters, alpha and beta, where this color here is 90% of the way there. This color here is 60% or 70% of the way there. But they're showing you that, for many things, it looks like 70% to 90% of the way there.

Here's their graph of profits versus the discount factor. So when discount factors are anything below 0.4 or 0.5, we're roughly getting 20% of the way to the monopoly profits, on average. But then, as people become more patient, the computers-- the more patient computers learn more and more to collude with each other. And when the computers are very patient, they're learning 90-something percent of the way to the monopoly price.

They have illustrations of what the price wars look like. So this was just an example. Take a final strategy. Someone undercuts. The other firm, then, does undercut them, and it takes them a while to get back to the monopoly price. So the computers have learned to follow some pattern like this, that has finite price wars. And this is their illustration of what one of the equilibrium strategies looks like, where this is the absorbing state where we charge the monopoly price.

And if you deviate, you follow one of these paths, like this, and then eventually come back, six or eight or whatever periods later. But the computers, in some sense, created this network of ways that you traverse this thing that eventually gets you back to the collusive price. And just, yeah, a cute illustration of what happens in these things. You do see a lot of stories on this of this happens on Amazon, that firms set these automatic pricing algorithms.

Then you go to buy something on Amazon, you see these crazy prices. This is just some illustration I found of this pillow was being sold on Amazon for $10,235. And clearly, it was-- there were two firms selling this particular Downright Eiderdown 434 Thread Count Silk Pillow, and the computers were playing against each other. And somehow, they went off into a part of the space where each one thought, if I'm a little more expensive than him, I'm a little more expensive than him, something great is going to happen.

And they ended up with-- they competed theirselves off to $10,000 pillow. And maybe the pillow doesn't-- once you get out here, you try undercutting from $10,000 to $9,000, it makes no difference. You don't notice your way back to a pillow should be $30 or whatever. But you can get lots of things that make it clear evidence that it does seem like, in a lot of products, there are algorithms competing, and the algorithms work sometimes and not others.

Any questions? So anyway, if that's it-- so next class, I'm doing empirical stuff. So it's Rob Porter's paper. If you want to pick one thing to read, Rob Porter's paper on railroad cartels. I have a paper building on Rob's paper and then, also, a paper about-- talk about a paper about Edgeworth price cycles in practice.