[SQUEAKING]

[RUSTLING]

[CLICKING]

GLENN ELLISON: OK. Let me go ahead and get started. So today, continuing my Monday theory, Wednesday empiric stuff, I'm going to talk about empirics of dynamic competition. In this set of lectures, I say there's also-- there's reduced form and structural literatures on dynamic competition, and here I'm going to completely ignore the structural literature. It's just there's some technical things that come at the start of that are techniques, and so we've left that all for [14.]273.

So those we do in [14.]273 in the spring, we'll see it there. But I would say that also it's one of the-- there's some neat structural stuff on dynamic competition, but it is a real challenge for structural estimation because a lot of structural estimation is built on there being a unique equilibrium and being able to invert going from the equilibrium back to the primitives.

If you think about everything I said about repeated games when there are a zillion different equilibria, if the same primitives can lead to a thousand different equilibria, you can't necessarily go backwards from, here's what we observed to what the primitives are, and that causes challenges for that literature. But anyway, I will not get in there.

So anyway, I'm going to talk about reduced form literature, which a lot of it is trying to provide evidence on, do the models that we've talked about previously apply and can we get anything about where they're working well and where they're not working well and what else we might want to put in them. I'm going to start with one of my favorite papers, Rob Porter's "Study of Cartel Stability: Joint Executive Committee, 1880-1886." You know, I have a number of classic papers in this course, a number of classic papers I omit from this syllabus.

Porter's is a rare paper that you pick it up and it's 40 years old and somehow, 40 years later, it still just reads very much like a current paper, seems very well done. It's very thoughtful and there's hardly anything in it when you look at it and say, wow, that seems archaic, the way he's doing things here. So anyway, paper is about railroad cartel in the 1880s. Again, if you go back to the 1880s, you have to remember a lot of things had not yet been invented in 1880.

The electric light bulb had not been invented. The car had not been invented. The paved road had not been invented. Obviously no need for paved roads if you hadn't invented the car. But so if you needed to transport goods, there were basically two ways to transport goods, by water or, then newly, by rail. And railroads became an incredibly important industry. If you look at the US stock market in the 1880s, I think the statistic is like railroads were roughly half the capitalization of the US stock market.

So in some sense, half the firms by value in this country were railroads. One thing that was challenging for railroads is a lot of this class has been about what makes price above marginal cost. Well, if you've got a railroad, you need to have price above marginal cost because they have these enormous fixed costs of laying tracks to get from New York City to Chicago and putting down steel rails and clearing space on the whole path and acquiring the right of way.

And then once you've done it, you can transport goods fairly cheaply. You can undercut the horse and cart by an enormous margin. But marginal cost of carrying the goods once you've built the railroad, are relatively low, and so if competition forces price down to marginal cost, you're never going to make the billions of dollars it cost you to build the rails themselves. So there were lots and lots of railroads who struggled to keep up prices. There were lots of railroad bankruptcies.

So Porter's paper is studying a cartel. So all of these things which are not so visible, but these-- I took the old maps. These are all railroads that went from Chicago out here on the shore of lake Michigan to the East Coast. In 1880, there were three ways to get from Chicago to the East Coast. There was the New York Central railroad. This is the one owned by Vanderbilt. It went either this way or this way along the Great Lakes, through Albany, and then down to New York City to get to the East Coast.

There was the Pennsylvania railroad that ran across Pennsylvania and got to Philadelphia and then went to places from Philadelphia, but got to the East Coast there. And then there was the Baltimore and Ohio railroad that started here, took a more southerly route across Pennsylvania, and ended up in Baltimore. And so from either of those three-- you could use either of those three railroads to get things to the East Coast. Jumping ahead a second. Why was this important?

Railroads did carry passengers, but passenger travel was not where most of their money came from. Most of their money came from transporting freight, in particular, transporting grain. Then, as now, this is where corn grows in the United States. Corn and other grains tend to grow in the Midwest. It's much cheaper to produce things there than in New Hampshire. So all the grain was grown out here. It got to Chicago, and then you needed to get it from Chicago to the East Coast, which is where a lot of people lived in big cities and didn't have farms.

And then in addition, a lot of the grain, when it reached the East Coast, actually went to Europe. So it would go by boat from the East Coast to Europe. And because it was going from the boat to Europe, in some sense, that equalized the prices at all the ports on the East Coast, because basically the European demand was big enough that roughly you didn't-- price getting to New York, Philadelphia, Baltimore was all roughly the same.

So anyway, cartel had three members. What did they do? So there have been a couple of failed cartels before that, but cartels were still legal until 1887. So what the cartels did is they hired the premier cartel consultant of the day, Albert Fink. Albert Fink met with them. You know, these firms could have merged, but there are many reasons why you may not want to merge. Each one had their different railroad. They wanted to do different things with it. So the three firms continued to operate separately and they just hired a cartel consultant to arrange how their cartel was going to work.

What do we know about what Albert Fink did? In some sense this was the secret to his business model, so he didn't publish exactly what it was that he did to maintain cartels. One thing we know he did with the railroad cartel is he had the option that observability was important, and he had the observation that observability is actually an easy thing to do in railroad because you can just hire minimum wage employees. Minimum wage employees were cheap in the 1880s. You could just hire minimum wage employees to count cars leaving Chicago.

And so you can just see, trains have New York Central or Pennsylvania Railroad or Baltimore Ohio painted on the outside of them. You can just count how many carts full of grain they have leaving the rail yard and you know the quantities from all the firms precisely. What you don't know is whether the firms are cheating on their agreements and charging less than $0.30 per hundred pounds or whatever they've all agreed should be the rate for getting grain from Chicago to New York.

So it's a situation where price is unobservable because you don't know if they've cut secret deals with people to carry their grain to the East Coast at a lower price. Quantity is clearly observable because you can just count cars. And those records that he has of his car counters survived, and so that's part of the basis of Porter's paper.

OK. One other detail is the 1880s was before a lot of global warming. So it was not 68 degrees in November in the 1880s. In Newton, where I lived, there were lots of big ice harvesting businesses where people needed ice boxes or whatever, and they would, on the lakes in Newton, they would carve big foot thick chunks of ice and sell them in the winter. Obviously that doesn't work anymore. One thing that also didn't work then was you could also get grain out of Chicago by shipping it over the northern tip of Michigan down here and then through canals connecting the other great lakes out to the Saint Lawrence seaway. This was the cheaper way to get grain out of Chicago.

There was a limited fleet of boats that would pass through the Great Lakes, and so in some sense, there's a fixed quantity that's going to go through the Great Lakes due to the capacity of the ships, but then that's fixed. It's always going to undercut the railroads, but it's just a limited quantity. And then importantly for this paper, pre-global warming, the Straits of Mackinaw were frozen for several months a year and were impassable for several months a year. So grains harvested in the fall, you want to get it out, but then in the winter, it's months in which the Straits of Mackinaw are closed and you can't get it out by ship.

Any questions on the basics of this situation? So anyway, Porter's thought was that interested-- there's previous literature talking about price wars in this industry and failed attempts to collude. Porter's is the hypothesis that, no, this wasn't an industry with failed attempts to collude. These price wars that people were talking about could have been signs that Fink had already figured out the Green-Porter theory of collusion under imperfect observability, and what we were observing is the functioning of a Green-Porter style cartel.

That's what he takes out to investigate. Starting point for his paper was then thinking about what would we observe if there was a functioning Green-Porter cartel in this industry. And so he starts with this assumption that demand takes a log linear form, log of qt, quantity demanded at-- and so these t's are going to be weeks in his in his paper. Log of quantity in a week as a constant minus alpha times the log of the price plus alpha 2 times the LAKES dummy plus u1t.

So this is a dummy for whether the Great Lakes are frozen, which he takes as just an exogenous shift in demand. OK. Let me say that's a better way to do it in a-- it would be better to do that in a linear demand model. In a linear demand model you would say there's some quantity, like 3,000 tons a week that disappears when the Great Lakes are open, and that we're subtracting. Why we're doing it additively in a log model is less clear, but anyway, that's what he what he's able to do in the log model.

And then he says, what's the cartel going to do? I know that our basic markup formula is price minus cost over price is minus theta over the elasticity. And he's going to suppose that a cartel is doing like what we said cartels do. They set high prices for a while. They jump down, they set low prices, they go up and they set high prices. They do this. And he has two different mark ups. He has theta equals theta c in cooperative periods and theta equals theta w in price war periods.

And you might think that if you're setting the monopoly price in the high periods, you do price minus cost over price equals one over the elasticity, so theta would be one. And in the price war phase, if you have go to pricing at cost, it's 0 over the elasticity so theta would be zero. That's not exactly what an optimal cartel does, though, because if you think about the first order condition, d pi dp is equal to zero when you're setting prices at the monopoly price. And so because d pi dp is zero, cutting prices slightly makes it easier to deter collusion and has no first order effect on your per period profits.

So in an optimal Green-Porter cartel, you always want theta somewhat less than one because it just makes it easier to deter colluding and doesn't do anything to your profits to first order. And again, the theta equals theta w during price wars. I said that simplest Green-Porter model you would want theta equals zero, you just price at cost. But a [INAUDIBLE] might note you could go to theta below cost or you could-- maybe you don't go all the way to cost anyway. But we think theta is going to-- this one may be a little bit less than one, this one may be about zero would be the natural assumption.

And then he says, well, OK, let me-- I've got a demand. I've assumed a demand curve. Let me also assume a functional form for marginal costs. He picks this one. So marginal cost is a constant times quantity to the beta. If there are increasing costs, you think beta is going to be slightly bigger than one. The more quantity we have to carry in a month, the higher is our marginal cost. So our quantity is being slightly higher than zero.

So beta equals zero would be constant marginal costs that are independent of quantity. Beta slightly greater than zero is marginal cost increased with quantity. And then a random shock in every period. The random shock could reflect fuel costs, labor costs, difficulty of getting-- your equipment not being in place to carry a load, needing to-- whatever. You may have different quality of equipment, and so you have some weeks you have to go to cheap to run equipment is available, some weeks it's not.

And so he says, what would happen if we had those marginal costs? Well, you would get this equation price minus marginal cost over price equals this. And so then you just solve that equation for price. I put price over to this side and I get price times theta over alpha 1. So I have price times 1 minus theta over alpha 1 equals this, and so I get price equals 1 minus theta over alpha 1 times that. OK, so this is how a firm would price, again, using different thetas in the two different periods.

OK. This may look like you've just got a big mess, but then take logs of both sides, and what you get is log of p on the left side. The right side, I get some terms that-- let me just do this first for your inner price war. So this is theta w. I first get some terms that I would get these terms, log gamma 0 minus log of 1 minus theta over alpha 1, and then I just get beta 1 log q. And then I get log of this is just u2.

And so then what Porter says is that-- so this is what I would get in a price war phase. I would get that if you had a price war. And then if you were in the cooperative phase, I'd get log gamma 0 minus log 1 minus theta c over alpha 1 plus beta 1 log q plus u2t. That's what you'd get if you were cooperating. And so what Porter notes is that this is just a linear demand curve where log of pt would be a constant, where the constant is this thing you get in the price war phase, plus beta 1 log qt plus u2t, plus another constant times an indicator for whether you're in the cooperative phase.

And the cooperative phase just shifts you from this to this by adding log 1 minus theta w over alpha 1 and subtracting 1 minus log theta c over alpha 1. So what he says is that if you had this-- again, two assumptions. You have this log linear form for demand and you have this particular form for marginal costs, then this is the appropriate supply equation. And it's going to also be linear. And so then the neat thing he's got is that basically this assumption of Green-Porter or cartel plus log linear demand gives us just a simple two by two supply demand system.

So demand is what he assumed it is. Supply is what's derived like that. And if you think about the way that system is working, I think Tobias-- or you've seen before, you're trying to estimate supply and demand. You need an instrument. You need a variable that's excluded from demand to estimate the demand curve and you need a variable excluded from supply to estimate the supply curve. And so what he gets is that we have a supply demand system.

The indicator for whether we're in a cooperative phase or a price war phase doesn't affect demand. This is a very common thing to do in IO. It is a great assumption. You're looking for what affects demand without affecting price, what affects pricing without affecting demand, whether the firm behavior determines their prices but it isn't directly affecting demand, other than through the price change. So if the firms are colluding or not colluding, that affect shifts price, but that doesn't have an independent effect on demand.

And then what does he have that has an independent effect on demand but not on price? The Great Lakes being frozen. So if you take this view that the Great Lakes are just pulling some amount of traffic out, but then once they've pulled it out, the remaining traffic is just the same residual demand you would have been facing before, it's not affecting your-- in this model, it's not affecting your pricing. So anyway, we can use i as the instrument for price in the demand equation, we can use lakes as the instrument for quantity in the supply equation, and estimate supply and demand.

Actually, as I said, this argument is better if you use linear demand curves. So this week's problem set, it's a largely empirical problem set if you haven't seen it yet. I'm giving you the data from Porter's actual paper and having you pretend you're Rob Porter and you think a linear demand specification would be better than a log linear demand specification, because in the linear demand specification, subtracting the LAKES thing is a better thing to do and it's asking you to redo his paper had you wanted to assume linear demand instead of the log demand. Or at least discuss how you would estimate the model.

OK. Anyway, so Porter's paper. Again, big data set for 1880 is a small data set for today. It's 328 weekly observations. What you've got in the data set here is total shipments of grain for the week, the official price. So Albert Fink himself announced what the price was. So the price is $0.30 per hundred pounds. That was part of his responsibility was just to tell the firms what price to charge every week. LAKES as a dummy equal to one if the Great Lakes were open for shipping.

So someone recorded the dates at which these lakes froze and unfroze every year. And then when railroads are half the economy, there were-- if you read the *Wall Street Journal*-- I don't know if the *Wall Street Journal* existed, but half of the articles would be on railroads because they were half of the-- half of the economy. And so there's plenty of newspapers-- there's actually a newspaper he uses called-- I think it's *Railway Review.* So besides being daily newspapers on business, there were daily newspapers on the railroad industry.

And so you could buy a newspaper called the *Railway Review* and read today's railroad news. Harvard Business School library has copies of *Railway Review.* I don't know if you can get in there with an MIT ID, I think not. But anyway, when I was a graduate student, I did go there at one point and look at the old railway reviews there in Baker Library. Anyway, he recorded whether there were stories in the *Railway Review* indicating that a price war is occurring in the industry, in the railroad industry.

And so if the newspaper had stories about price wars going on in the railroad industry, then he would report it as being one, otherwise he would record it as being zero. And so in some sense he used this as his cooperative phase is no story in *Railway Review* about price war. Price war phase is, *Railway Review* says price war going on in the JEC cartel this week.

So we get two ways to estimate this model. One, the simplest one, is just assume that i is observe-- the cooperation indicator's observable and just use the PO variable that he got from the newspaper to do this. You then just have a standard simultaneous equation model. The other approach that he uses is to assume that It is unobserved. So it's an unobserved random variable with some mean gamma, which is to be estimated, and then you have a simultaneous version of a switching regressions model.

Have switching regressions model entirely disappeared from econometric sequence? Yeah? OK. So OK, how does this switching regressions model work? So suppose you have the-- obviously our standard linear model, like y equals x beta plus epsilon, I put x here. There's some relationship Like y equals x times beta. And what you expect to see is a cloud of points like that. OK? So you've got y equals a function of x. That's what you observe. You can fit a line by OLS going through that.

We do that all the time. Suppose you have the following. Suppose there are two different relationships. There's y equals x beta plus epsilon, and there's also another relationship, y equals x beta prime plus epsilon. And this happens with probability 1 minus gamma, and this happens with probability gamma. OK?

So suppose there are actually two different relationships. For instance, if you were doing demand for ice cream on a beach, there are two different relationships, demand when it's rainy and demand when it's not rainy. You may not have raininess, but you know that there's going to be one demand curve on rainy days, one demand curve on non rainy days, and they're going to be different.

And it may be that this is the demand curve on the non-rainy days as a function of price-- I guess the price should be positive-- and then on the sunny hot days, there's some other relationship that looks like this. and you might ask yourself, OK, suppose it is just-- suppose there's one x, but then there's also a constant, and suppose that it really is just this is the y equals x prime beta plus epsilon, and the constant is higher on the sunny days and demand looks like that.

OK. Question is, can you estimate that? OK? First answer is without any additional assumptions on-- if you just go with a standard expectation of epsilon given x equals zero assumption, you can't estimate that. And the reason you can't estimate that is you might think, well, can't I just draw two best fit lines and ask what combination of two lines minimizes the sum of squared residuals from the closer line, and isn't that an efficient estimator of this model?

But the problem is you can't tell this model apart from a model where there's one line and the epsilons have this distribution. If you're only doing mean zero errors, you can't tell two lines with this distribution of errors apart from one line with this distribution of errors. And so without some assumption beyond expectation of epsilon given x equals zero, you can't estimate the switching regressions model. But if you're willing to go epsilon normal zero sigma squared.

If you're willing to add that to your model as an identifying assumption, if you're willing to add this, then yeah, the simple intuition that I can estimate this one by maximum likelihood and just saying what combination of two lines and sigma. So what beta, beta prime, sigma, and lambda best fit the data, I can do that. So you need functional forms to estimate switching regressions, but if you're willing to impose functional forms, then you can estimate this and you can recover the-- you don't ex-post know which equation generated every data point.

So there's a data point right here. You don't know for sure whether it was this line with an error of zero or this line with a very huge positive error, but you do get a Bayesian posterior that, with high probability, it was this line that generated it. So what you get is you get this model, and then ex-post, you get a probabilistic Bayesian posterior on which state generated every data point. OK? Anyway, so what Porter does is just basically do the simultaneous equation version of a switching regressions model.

And he says that we have these two models that generated the data. Every period is either generated by this with I equals 0 or this with I equals 1, but I is unobserved. And all I know is that there's probability-- there's some probability lambda that I equals 1 or that It equals 1. But I'm going to treat It as an unobservable and fit the data by a combination of two simultaneous equation systems, one simultaneous equation system with I equals zero, one simultaneous equation system with I equals one.

And in every period explain-- again, it's every period you get a ptqt pair, you just ask which model best explains ptqt, and is it the model I equals zero or I equals one, but I'm going to do this by maximum likelihood. I think that's, in some sense, the preferred estimation in his model is to say that I don't want to trust the *Railway Review* data and I'm going to do this with the unobserved It's and say what fits the data better.

I think this idea of explicit unobservables and putting the unobservables in your model and making some assumption about the random process they follow, huge. This has also been something that's done all over the place in IO these days of taking an unobservable, putting it in there, modeling the unobservable, and then seeing what you can do with that unobservable.

OK. What do his estimates look like? The first thing is that prices are dramatically higher in the cooperative phase than the non-cooperative phase. If you use the PO dummy, prices are e to the 0.38 times as high when you are in the cooperative phase. When you use the PN dummy, prices are e to the 0.545 times as high. So what is e to the 0.545, like 80% or something? So prices are much, much higher. Here's a simple table.

Again, we don't know PN, but we've got PN star, which is what's our-- classify them ex-post based on which do I think it is after seeing the price. And it's like the average price is 17in in the cooperative phase, 28 in the non-cooperative phase. Sorry. 28 in the cooperative, 16 in the price war. So that's roughly the magnitude of how much collusion is affecting prices. The model with the switching regressions, assuming that the PO is misreported, fits dramatically better than the one using PO. The r squared goes from 0.32 to 0.86.

So it looks like the model with the unobserved cooperative phase is really important. The model tells you what percentage of weeks are in the price war phase, and it says that about 28% of the weeks are in the price war phase in the data. And if you just look at the raw data, I guess the beginning of the first price war, you see prices going on this at $0.40, 36, 35, $0.30 per hundred pounds, 35, 35. And then it drops to 16.

Again, this is not some cost shock that you suddenly became able to transport grain from Chicago to East Coast at half the price as of one week ago. It does look like just dramatic evidence that this is a price war. They price really low for a long time. We do think they would be losing money getting grain to the US, the East Coast at $0.12, and then they jump back up afterwards.

I didn't put that other key test in. Obviously, the key test he asks of the Green-Porter model is, does a model with two supply curves fit better than a model with one supply curve? The static Nash story here would be that we have this model and there is no I. The I is just always-- they're just always behaving the same in every period. And so if you test these two different behaviors versus one behavior model, I guess you can think of that as like testing lambda is not in either zero or one.

Lambda equals zero or lambda equals one is the model with one regime instead of two. And can you tell whether lambda is significantly different from zero or one? I don't have an estimate of lambda in here, do I? OK, anyway-- He provides tests in the paper and saying it's tremendously significant that two supply curves fit better than one supply curve. You know, it's not surprising. It's things like this are really hard to fit with one supply curve, whereas with two supply curves, you fit that much better.

OK. So next I get to go to one of my own papers. Still including this, just sentimental reasons. So this was actually when I was a graduate student here. There was no second year paper. So 14.192 was not invented until the year 2019 or whatever. What did exist before then was there was both a history paper and an econometrics paper.

And one of the great bonuses was for-- what everyone tried to think of as getting through the program with the least amount of work possible is could I write a history paper that's also my econometrics paper? Because you then got to-- there was no rule against handing in the same paper for your history paper and your econometrics paper. So as a student in this program, I-- oh, this railroad cartel data is from 1880. That qualifies as making it a history paper.

So anyway, so this was my history paper and my econometrics paper. But anyway, and I guess I also-- just as combinations of you don't have to be all that creative and original to write papers. And a lot of the time it works to just make small improvements on the existing literature and just-- you start with something that's in the existing literature. You make a small improvement of it. You know you're at the frontier. Whereas if you start from scratch and do something really creative, you can end up with a worse version of what someone else did before.

And so anyway, so this was my paper looking at-- thinking about what more could I say about Green-Porter. And the motivating idea here was that Porter shows that there are two regimes, and two regimes explain the data better than one. But I think the content of the Green-Porter theory is not just that, but price wars occur when there are suspicious demand realizations and demand realizations that look like someone cheated on the cartel agreement, and then when there is a demand realization that looks like someone cheated on the cartel agreement, then you should be more likely to trigger a price war.

And what we think these price wars should look like is you're pricing high, you go down, you have a price war for some number of periods, and then you go back to pricing high again. So it's both the price wars are supposed to be continuous periods and they're supposed to come after some demand pattern that looks like someone might have cheated. And so that's what I'm trying to test here. So basic exercise is just re-estimate Porter's model. I do make a few changes which reflect this being the 1990s instead of the 1980s and computers were better.

One is paper allows for serial correlation in demand. The demand in Porter is tremendously serially correlated if you look at it, so he didn't allow serial correlation. Supply is just like in Porter, with there being an unobservable It. But then the main economic idea is to let the regime transitions follow a first order Markov process, where the probability of cooperation at t plus 1, conditional on whether you are cooperating at t and some other observable variables wt, is going to be e to the gamma wt over 1 plus e to the gamma wt.

So these variables wt are going to affect the probability of the Markov transitions from cooperating to price war or from price war back to cooperation. And what are the variables I put in here? So obviously first version to estimate is just let w be constant. This is Porter's model where the regimes are i.i.d. If there is nothing in w other than a constant, this is Porter's model. If you put It in here-- so you get the cooperation indicator affect-- you let It affect the probability that It plus 1 is one.

Then you have first order Markov structure on price wars, and you can see whether price wars are continuous. Because if price wars are continuous periods, what you'll see is if you're cooperating, there's a 90% chance you're going to cooperate. Once you're in a price war, there's a 90% chance you're going to be in a price war, so the price wars become long, continuous periods. And then the primary interest in specifications that also include variables that look like suspicious demand patterns.

What is it that Fink was doing? Well, Fink could have done many, many things. There are lots of things that could make you think that somebody was cheating or could make the firms think someone was cheating. One would be whether any firm had an unusually large market share. This isn't in the Green-Porter model, but in the JEC, there were guys counting cars and counting cars by railroad. And so if you observe the Pennsylvania Railroad has 50% more demand than everybody else this week, it could be that all the grain shippers were actually trying to get it to Philadelphia, but it could be the Pennsylvania railroad was cheating.

So if one firm had an unusually large market share, that might make you trigger a price war. Another thing would be if any firm had an unusually small market share. Firms could worry that the other two guys are both cheating against me. And so if I see that instead of being 30/30/30 or 30/30/40 or whatever, it's 40/50/10, the firm that's getting the 10 could worry that the other two are both cheating on the cartel agreement. And so you might want to have a system where, if anybody has an unusually small market share, then you start a price war.

Another thing you could also do is just if aggregate demand is unusually high. This would be more like the Green-Porter textbook model, but if suddenly aggregate demand is really high, then you think that someone out there must be cheating, therefore we're going to be more likely to start a price war. And obviously you could do many combinations of these. What I do in the paper is just put in several of them and see what's there. First is the standard model without the w's.

So standard model without the w's, there's strong serial correlation. The effect of collusion, or the cooperative phase, on pricing is even larger than Porter by a little bit. It's 0.6 instead of 0.5. But what does change when you account for the serial correlation is the demand elasticity, which Porter had at 0.8, and it was a puzzle. Why would the cartel be pricing where the demand elasticity was 0.8? Don't we know that you normally price in the elastic part of the demand curve?

Once you account for the serial correlation, that mystery goes away and I now say that the demand elasticity was minus 1.8, not minus 0.8. And so if you think about what the implied theta by the data is, the implied theta from how they're pricing-- again, you have price minus cost over price equals theta over the elasticity. So now you plug in 1.8 here instead of 0.8 and you say, what's the theta that's causing this? Now I'm getting that the theta is fairly-- they're using a theta that's pretty close to 1. It looks like the theta is like 0.85.

So again, you're not-- if you assume that theta c equals zero would imply that theta-- sorry. Theta w equals zero would imply that theta c is about 0.85. So with this new demand elasticity, it looks like it's not just a Green-Porter cartel having two regimes. It's two regimes where one is close to cost and the other is very close to the monopoly price. So it looks like they are aggressively trying to get nearly the full monopoly profits. And if you look at this, the probability of a price war-- the probability of cooperating when you're cooperating today is e to the 3.6 over 1 plus e to the 3.6.

So if you're cooperating today, the probability of cooperating tomorrow is very high. If you're in a price war today, the probability of cooperating tomorrow is e to the minus 2.6 over 1 plus e to the minus 2.6. So that says if you're in a price war, there's very low probability that you're going to go back to cooperating next period. So clearly that these numbers are different mean that these price wars are continuous periods in the data. Also it makes you reclassify some of the periods.

In Porter's data, you can be like, OK, price war, price war, price war, price war, cooperation, price war, price war, price war. And then with putting in this variable that reclassifies them and says, no, that's a price war too. It's just a price war that has a big epsilon attached to it. OK. OK. And then what about these variables for things causing price wars? Some of them are kind of significant if you do at 5% level significant one-sided tail.

So some firm having an unusually large market share, it seems like that makes price wars more likely to start. Also, VT, the aggregate demand residual being high, also makes it look like price wars are going to start. So if think about this is a model of Albert Fink, it does seem like if aggregate demand is unusually high or one firm has an unusually high market share, that makes it more likely that he's declaring price war starts tomorrow.

Significance is just marginal, though, and the triggers are not really powerful enough to deter deviation. So if these were the magnitudes of the triggers that Fink was using and you were advising the New York Central Railroad, you'd be like, yeah, you guys should be cheating. Because if you cheat, you gain 10% market share, you're going to earn a lot more money, and the probability of price war only goes up by a few percent so you'd be better off cheating on this agreement. So obviously, what's going on is-- this leaves multiple hypotheses.

So one is that Fink got it wrong and his triggers weren't sensitive enough. He was too hesitant to trigger price wars, and therefore, firms did have an incentive to cheat and firms did cheat. The other hypothesis is that he had a trigger that's imperfectly correlated with all the variables I've put in there, and because it's imperfectly correlated, there's attenuation bias, there's a stronger trigger out there, I just didn't find it. Unfortunately, this is the kind of thing where you might say, could I do some machine learning thing and figure out exactly what Albert Fink was doing?

There are between 8 and 11 price wars. So in some sense, you only have 8 or 11 data points. There's only so much you can do with sophisticated statistics when you're trying to figure out when a price war started. OK. Any questions on that? Anyway, so other things in the paper. I do have something looking at Rotemberg-Saloner effects. I won't talk about that. The paper asks the question in the back part, are there secret price cuts, and does it look like secret price cuts were being given?

And the way it does this is to add extra unobserved demand regimes, saying that if firms are giving secret price cuts, demand is going to look unusually high relative to what demand normally looks like. Can we estimate a hidden demand-- a hidden regime model to say, does it look like there's some fraction of periods in which people are cheating and demand is unusually high? And it provides some suggestive evidence that it does look like that's probably what was going on here historically is that there periods where demand is suspiciously high, making you think someone is cheating.

And so we think the-- it aims toward that Fink didn't have it quite right and the firms were doing some cheating on each other. But obviously, if you think about where did I miss an enormous opportunity, there are papers written subsequent to mine that now have thousands and thousands of citations that say, well, what you could do in a situation like this is instead of estimating the trigger and saying, is the trigger strong enough to cause collusion?

Maybe what you could do is impose that the trigger is strong enough to stop firms from colluding, and then estimate the model and see if you can find the trigger based on that, allowing some extra measurement error. So it could be that the price wars start one week earlier than I think they start because some firm had an enormous market share one week earlier. The reported price is still high, but maybe the price war actually started one week earlier with a mismeasured-- there was an enormous cheating incident for two weeks earlier. They start the price war. The price war only shows up in my data, which uses posted prices one week later, therefore I miss it.

So anyway, I think this is, in some sense, what the structural literature has done is to say, if we think that the firms are playing an equilibrium, we should impose that an equilibrium is being played and make that an extra constraint on the model and estimate the model with that constraint. I didn't think to do that at the time, but in some sense, it is a very natural thing to say is let's impose that constraint, see what happens, and is the model now able to explain the data better.

Yeah. So for instance, in this week, it's very clear where this price war starts. The price war starts here. But it could be that in this week the firms were already cheating on the agreement and they were actually pricing here, and we just have a mismeasured price because the prices in this data set are the prices that they were supposed to be charging, not the ones they actually charged. Next thing I want to get to is this is something I had the benefit of seeing with my career. It's interesting.

I read *Mass Control* for many years, and *Mass Control* would point out that, be like, these Edgeworth cycles are possible. It was always thought of, at least seen as this is an interesting possibility that, in models with small numbers of firms, firms can collude and they could collude not on constant prices like in the standard repeated game equilibrium, but they might collude-- collusion might lead to these Edgeworth cycles where price goes like that.

And I think it was always at least initially thought of as a intellectual curiosity that collusive models might actually more naturally lead to cycling prices rather than high prices. And then this is a view that the world actually has changed as high frequency price data came into-- started getting collected because, again, if I ask you, like, what are the prices at gas stations, and what do prices at gas stations do? You know, you may know that like, yeah, sometimes gas is 3.69, sometimes it's 3.79, sometimes it's 3.59.

I've never really noticed a pattern other than knowing that sometimes I get a good deal on gas and sometimes I get a bad deal on gas. And so actually, Mike Noel was a student sitting in this class where you guys are sitting in the early 2000s, and he decided he would actually investigate the Edgeworth cycle hypothesis. At least had some notion that it seemed like-- he had the view that in Toronto, where I grew up and went to college, I think prices did cycle, or at least I know they changed a lot and I wonder if they are cycling.

And so what Mike did was he still had a girlfriend who was in Toronto when he was a PhD student, and he got her to collect a data set for him. And the way she did this was she had a, like, 25 or 30 minute car commute in Toronto, and she stuck a voice recorder in her pocket or whatever. And as she passed gas stations on the way to and from work, she just read to herself what prices they were charging, and then transcribed them.

And so what she got was a twice weekly data set of prices of gas stations in Toronto which are located along-- it was like map of Toronto, she works here, she lives here, she drives this route every single day. And so you get 14 gas stations along that route and you get the twice daily prices. And this is one of those ones where just the immediate paper, this is like-- so this is a graph of her twice daily data.

And what you see is that this was the wholesale price of gasoline in Toronto if you just bought it wholesale at whatever those big enormous fuel tanks are that the pump trucks back up to take it to the gas stations. And this is her data on the retail price of gasoline. And I just think just that one graph just makes it enormously clear that this is not a cost-based explanation for why gas prices change. Clearly, these guys are playing-- somehow the gas stations in Toronto are somehow playing some kind of Edgeworth cycle equilibrium, and they're playing some kind of Edgeworth cycle equilibrium where the prices just jump up really high and then go down.

The firms at various points are really losing money, pumping gas, and then the prices jump up and become very large, and then they just keep cycling like that. These cycles, you might think, oh, are those weekly cycles? The frequencies there are roughly a week long, but they're not synchronized with days of the week. Sometimes they're six days long, sometimes they're eight days long. Over the course of the day, the cycles shift relative to the week. So it's some weeks the big jump comes on Mondays, some it's Wednesday, some parts of the day it's Friday.

But anyway, prices are just following that cycling pattern. And actually, I find this kind of remarkable, because if you think about it, somehow, there's a interesting network geography question here of, how do they even do this? I get that, OK, this station raises its price and so then the two next to it raise their prices and undercut it the next day, and then this one undercuts this, and then these two undercut these ones and then the even ones undercut the odd ones or whatever.

But somehow there are lots of chains of gas stations. You could also have driven this way and you would have gas stations here. How do you keep the cycles in sync so that-- you might think like, I go this way and I pass an even number of stations. I go this way, I pass an odd number of stations. Who's undercutting whom? But somehow the whole city geography must be such that the city hangs together and cycles together with itself.

And actually, someone has another paper on Canada more generally about high frequency data, and it turned out that-- and I don't know why-- some cities in Canada do this and some do not. Some cities have these cycling prices, some do not have these cycling prices. And again, it was largely unknown until people started gathering high frequency data, because you use data at a one week frequency, you don't see this because you miss the periods in between. Another subsequent paper on this, Zhongmin Wang wrote a paper about gasoline pricing and-- yeah?

AUDIENCE: If we believe like that foresight story, wouldn't we believe that these cycles would only occur with gas stations that are close to other gas stations? So if you have gas stations out away from other stations, they would have no incentive to do this.

GLENN ELLISON: To do this. Yeah. So obviously, this is-- yeah. This data is urban Toronto. Toronto is a city of I don't know how many million people, so I assume that there are gas stations every-- I assume every gas station in urban Toronto is close to some other gas station in urban Toronto. And so it could be something about, yeah, the density needs to be high enough to get the things to cycle, and once you get islands without gas stations, the islands break up the cycling in the other places.

Yeah. You need to think of there being price competition between adjacent stations and every station's adjacent to another one to have this to work somehow. I think these graphs are the average prices in her data set across all 14 stations or something like that, so it is like a synchronous-- at least in her commute, it's synchronous for the length of her commute.

OK. So anyway, Zhongmin Wang wrote a paper then about gasoline cycling in Perth, Australia. Perth is also nice because Perth is all the way over on the West Coast of Australia where there's almost nothing else, and so Perth is like an isolated city all by itself. There's also only two-- there's actually one gasoline refinery in Perth, so everyone has-- you can import gasoline, but other than other than importing gasoline, you're all buying it from the other refinery, I believe.

Anyway, and Perth had this problem of prices there were known to cycle. And prices were known to cycle, and this was seen as something that was very bad for consumers. And so the Australian-- I don't know, whatever it is. Australian Competition Authority or somebody in 2001 passed a regulation to try to stop price cycling. And what they did to try to stop price cycling was a law, and the law said that gasoline stations can only change their price once a day at most.

Every gasoline station must, by 2:00 PM every day, report to the government what price they will charge the next day, and they must charge that price the next day no matter what. So every day by 2:00 PM you had to submit to a government website what price you were going to charge the next day, and then your price was fixed. And they were thinking that, OK, this is going to stop the-- this could stop the price cycling because in these price cycles, you go from pricing low to pricing $0.20 above everybody else.

If you do that, no one is going to buy from you. And then if we do this daily thing, firms are going to be left with zero demand for a full day. They're not going to want to do that. This is going to disrupt this price cycling which we think is keeping prices above marginal cost. Anyway, that was what they did. So anyway, and what Wang does is gather data before and after the gas price change. He had another very clever thing in his paper, which is how do I get high frequency data on the prices that gas stations are charging?

Can I go to every gas station in the city of Perth and say, could you give me your price change with the exact minute by minute time at which you changed your data, you change your prices? And the answer is obviously, no, you can't do that. The stations won't give you the price. He didn't have a girlfriend in Perth. You couldn't drive around to all the gas stations, and how frequently can you do that anyway? But what he did is he made a deal with a credit card company and realized that credit card companies have people dip their credit cards at all the gas stations.

And so the credit card companies know what the price was because it actually gets reported back to the credit card company that they bought 8.367 gallons or whatever, liters, at 12.99 a liter. And so he made a deal with one credit card company to give him gasoline prices in Perth. Every time that they recorded a swipe in one of their cards. You don't need that big of a market share in credit cards to be basically getting most of the gas stations pretty often.

Anyway, so we have these, again, very neat figures. So anyway, did I say these things already? Yep. Edgeworth cycles were present. Regulatory change. Yeah, daily prices for after the change you could get from the government website, but he got data from a credit card company for six months before. Oh, and wholesale prices you could also get because the single refinery that refined gasoline-- how do you have a market with a single refinery?

What the single refinery did is it committed to what price they would charge to not exploit people, which was the price of gasoline in Singapore plus 12.6 cents a liter or something like that. By observing the wholesale price in Singapore, you actually know the wholesale price in Perth. Anyway, so here we go. This is, again, the pre-data that he got from the credit card company. This is what gasoline prices did in Perth, Australia before the law. And again, you can see there's just this striking pattern of low price, jump up, charge high prices, prices drop down. These things are a week apart.

So the cycles are generally about a week long, but you can see they're not exactly week long. This one here, this one looks like it's almost nine days for that cycle. This cycle is shorter. That one's, like, six days. And you get these just massive one day increases in gasoline prices. These prices are-- these are cents per liter. So if you can't see it, this is 5. So that's, like, $0.30 a gallon of gas. So it's like the price goes from 3.39 a gallon to 3.69 a gallon in the space of an hour or a few hours, and then it spends a week going back down to 3.39 and then jumps back up to 3.69 again.

OK. Anyway, so what does this paper say? OK, they do this. They put in the law. What happens? I didn't put in all the graphs, but shortly after the law is passed, you get a bunch of chaos. The price cycles disappear. The price cycles disappear. The firms look like they don't quite know what they're doing. They're finding a new repeated game strategy. And then within three or four months after the law has gone into effect, boom, there's the daily price cycle data. So whatever learning of how to play a repeated game strategy and how to match these things up in this new environment where you can only change your prices once a day and are stuck with whatever you said you were pricing at 2:00 PM, price cycles return.

This period, he's got the wholesale price data in there, and you can see at some times they get fairly close to the wholesale price. And again, I don't know what the-- I'm not sure if this wholesale price includes shipment to the gas-- trucking fuel from the wholesale terminal to the gas station. If it doesn't, then these prices could be falling very close to their actual delivered price of the gasoline.

But what you are getting is this clearly more than doubling of margins where you have a periods of time when you have a very small margin, which could even be zero depending on what the delivery charges are, and then prices jump up very high and then they drift back down. So whatever learning or whatever intuiting of repeated game strategies does not stop gasoline prices from cycling here. What he also says-- he just provides other observations that might help you think about where this cycling comes from.

One is that before the regulation, these cycles had this very distinctive pattern where it was BP who was always leading the cycle. And there'd be some time between 11:00 AM and 2:00 PM on a Tuesday, Wednesday, or Thursday where BP would raise its price by $0.30 a gallon. And then Caltex, which was the biggest firm, would follow BP about two or three hours. Within two or three hours of BP raising the prices at all their stations, Caltex would bump all theirs back up to where BP's was or slightly under BP's price.

And then another hour later, Shell would be up there joining them. And then the others, the other small independent stations, might take a day or two to catch up and notice all the others had gone up to the high prices and they'd do it. After this pattern, something that seems like it disrupted this was that it seemed like BP was no longer willing to do this. When BP could raise its price and then have Caltex follow two hours later, it seemed like it was happy to raise its price and sell very little gas for a few hours.

It seemed unwilling to do this when it would be selling nothing for a full day and it didn't know-- maybe two days. Maybe Caltex isn't going to follow you the next day. Anyway, the pattern of leadership changes, and so now it seems like BP and Caltex share the role of being the price leaders. So maybe 40% of the price cycles are triggered by BP raising prices, 40% are triggered by Caltex raising prices, and also Shell is sometimes doing it, but again, doing it less often than the others.

He argues that it does look like-- you remember I talked about the way these cycles work is you go down here, and then at the bottom in the mass control model, you have this mixed strategy equilibrium where the firms mix between staying here and getting a low profit and going here and getting no profit. He says a model with independent mixing by BP and Caltex does look like it explains what's going on. It's like they're each mixing, and with probability 0.5, they go up, and with probability 0.5, they stay down.

He argues that it looks like independent mixing because you-- imagine if they were mixing 50/50 independently of each other, what you should see is a quarter of the days nobody goes up, a quarter of the days, one of the two goes up. A quarter of the days, BP and not Caltex. A quarter of the days, Caltex, not BP. And then a quarter of the time, they both go up. Looking at the frequency with which they both go up and none go up and one but not the other, it does look like they're playing independent mixtures down at the bottom, and that's the thing they've settled themselves into.

And basically, if you want, did this law affect average markups over wholesale price? This law did not. So whatever you thought that this was an anti-market power law, this did not work. And you know, it did not work despite the fact also-- I guess another part of the law was that you have to report your price, and those prices are being posted by the Australian government on the internet so that anyone wanting to buy gas now has the perfect gas app because everyone's required by law to charge the same price they announced at 2 o'clock yesterday.

So you can just get on your app in the evening and say, these are tomorrow's prices for every gas station in the country or these are the current prices of every gas station in the country. In the city, these are what they're going to be. Plan your trip to go to the cheapest gas stations, cause more price sensitivity. That didn't happen either.

Any questions? And, you know, in some cases, maybe why shouldn't it happen? Mass control is with perfect competition. With infinitely elastic consumers, you get these cycles. You make people closer to the mass control ideal of infinitely elastic, you still get the undercutting process. So it's not at all clear that that ought-- again, in a static sense, yes. In a static model, make people more price sensitive, prices go down. But make you more price sensitive, maybe it just makes the mass control thing even more of an equilibrium than it already was.

AUDIENCE: [INAUDIBLE]

GLENN ELLISON: Yeah, so I haven't seen-- you know, I don't know. You guys probably don't go. Leon Musolff was giving the-- he was giving the Harvard MIT IO workshop talk. You know, he was just doing data on pricing of goods sold on Amazon. So he has very high frequency data of prices sold on Amazon. It sounded like he was saying in his Amazon data it looked like there were Edgeworth cycles.

And he has an Edgeworth cycle paper that he's writing on it, but he was describing it as maybe it looks like 10% to 20% of the goods being sold on Amazon he's finding Edgeworth price cycles, and the other 80% he's finding things that look more like stable pricing. This is like people selling umbrellas on Amazon or people selling whatever. He does a lot of fashion categories. So anyway, I think-- I don't know. I haven't seen papers about high frequency data on rental car prices or airline prices or whatever to ask how many of those things look like they have short run cycles versus how many of them don't look like they don't have short run cycles.

But something that does go-- in mass control, it is, in some sense-- I think mass control does say you need to have pretty price sensitive things to create these cycles. It might be the internet price comparison sites make more things like gas stations that used to be like gas stations and might cause more of these than we used to see.

OK. So Then the final thing I wanted to do is give a little bit of-- again, this is an empirical part of this paper. This paper, really, it's an empirical part motivating a theory model, but I figured I would just discuss the descriptive evidence as what evidence do we have on pricing algorithms and online sites and what are pricing algorithms doing to pricing. And actually, it's more suggesting that we need different models of pricing algorithms. But so what Brown and MacKay do, again, this is a totally graduate student capable paper where they've got no proprietary data at all. They're just web scraping prices.

And what they did is they set up scraping programs to collect data on an hourly basis from April 2018 through October 2019. So it's about a year and a half of data. And every hour they went to five different websites. They don't give the names, but they do say these are the five most common websites in which people buy allergy medicines in the United States. So I assume the five websites are Amazon, Walmart, Target, CVS, and Walgreens, but I don't know. Anyway, they picked seven allergy drugs, Allegra, Benadryl, Claritin, et cetera.

And they just scraped the prices for every different package size of those items. Because you can buy Xyzal in 80 tablet jars and 55 tablet jars and 30 tablet jars and something or other, these seven drugs create, like, 150 different products. There are a lot of different packages. Many of the stores only sell the three most common packages for them. But anyway, so what they scraped is every price that they could find for any version of these drugs every hour for a year and a half.

And then they're trying to record basically, what happens to learn about the pricing algorithms that these firms are doing. Thinking about the q-learning theory. Are these firms running q-learning algorithms? What are they doing? What can we pick up from just observing the price changes that they're making as a function of their rivals prices and so on? There's clear heterogeneity in practices across website. There's one retailer they call A. I'm going to guess that that's Amazon.

Retailer A is very different from the others. One thing that retailer A is, is clearly is much bigger than the others. Every other website has roughly 50 of these drugs available. Retailer A has 150, so retailer A just has many, many more versions. Retailer A seems to have just about everything in the world that's sold is on retailer A, whereas it's not on the others. Retailer A is also changing its prices 1.89 times per day per product.

So if you ask, do you want to buy 80 tablets of Xyzal, how long do you have to wait for the price of Xyzal to change? The answer is, like, 15 hours the price is going to change, and it just changes over and over again. So when they graph the raw price data, the black is the retailer A's prices. Yeah?

AUDIENCE: Anecdotally, I know when you-- I think I've heard that sometimes when you look at Amazon a lot and keep putting, they just change the price for you.

GLENN ELLISON: For you personally.

AUDIENCE: So could it just be that Amazon's doing personalized pricing.

GLENN ELLISON: So it could be they're getting something that Amazon is picking up their program. And I don't know-- I don't know the details of what kind of program they were using to access this data and whether Amazon could have recognized that program and done something to it. But anyway, this black is the graph of what Amazon was returning to them, which is these prices changing quite frequently. Other thing you actually notice is price levels differ.

It looks like this A was changing prices all the time and substantially undercutting what the red and the green and the blue firm are doing. OK? The other one that changes price fairly frequently is the green firm, B. B changes prices, on average, 0.28 times per day. And so firm B is this green line here that's doing that there. At the end here it's harder to see because the green is more mixed in with the black. And then you have the three firms, C, D, and E, that are the ones up here. E is really like this blue line that's largely here. Red and the orange ones are going up and down occasionally.

And there are gaps in this data because you write a program to scrape data every hour and then the websites change their interfaces and you don't notice it for a while that your program is broken, and then you have to fix your program. Many, many times during a year and a half, they changed their website in a way that your scraping program doesn't work and you have to fix it, and then you have missing data.

So some findings. Firms seem to be updating on a schedule and there seem to be two technologies going on here. So remember, this is 2018, 2019. Firms A and B seem to be running some kind of price updating program that's roughly equally likely to change the price at any hour of the day or night on any day of the week. So these are the probability that a price change occurs in a particular hour of the week, running from Saturday midnight to Friday, 11:59 PM or whatever. OK?

A and B seem to be having these programs that just run all the time that change their prices. Retailer C seems to run a program between 3:00 and 6:00 AM and sometimes change prices. Again, remember, retailer C is changing prices roughly-- each individual drug's price changes roughly once a month at retailer C. So they seem to change prices roughly once a month. When they do change prices, it's between 3:00 and 6:00 AM on some day of the week, and all days they're doing price changing activity.

And retailers D and E are changing their prices at midnight Eastern time on Sunday, going into-- Saturday night, going into Sunday morning. So slightly after midnight Saturday night, they do change their prices. Retailer E, it's almost always right around midnight, and then retailer D, it's between midnight and 3:00 AM or something the prices change. So clearly these firms are using different technologies, and it seems like two of them are using a very old fashioned technology of having a meeting where they get together and discuss what prices should be, and then others are doing the change it more often.

Are these pricing algorithms like the pricing algorithms I discussed in the q-learning theory where your price depends on the other firm's prices and you're reacting to them? And yes, firms A and B, their algorithms are reacting to prices. What they do here is that-- you remember D makes those changes very, very rarely. D is making those changes once every two or three months, and it's always making those changes just after midnight. And so they ask, do A and B do something different on those rare Sundays when D changes its price?

And what you find here is that, yes, A and B are reacting to D. But A and B are not reacting all that quickly to D. So it's not like they're doing hourly reactions. So this is, is Amazon reacting to what goes on in cvs.com? And the answer is yes, but not for a while. So if you look at the probability that there will be a price change after D changes its price, and the answer is for 24 to 36 hours, there's no additional probability that A has changed its price based on D having changed its price.

But then starting in about 36 hours, going up to 48 or 60 hours, yes, they are then reacting to what D does. And there's more probability that they will change their price in the 36 to 72 hour window after D has changed its price than there would be otherwise. And so again, 36 to 72 hours later means looking, it's Saturday midnight. D changes its price. And so then it's Monday at noon to Tuesday at noon. A is more likely to change its price if D has made a price change in that period.

So it looks like A is running a price responsive algorithm, but it's not something that would cause the mass control gasoline price cycles thing that CVS changes its price, they're immediately undercutting. It seems like a slower undercut. But if you think about that, that might make sense because if you know that D isn't going to change its price for a month, you can just-- or two months, you can figure out, what do we want to do now that D is cheaper than us? We'll just pick a price and then we'll stick-- and then we'll be optimally responding to them. We don't need to respond to high frequency because d isn't changing.

Also what they find in this paper is that the frequency of price changes is clearly related to the prices being charged. So if you make it price index where A's price for a product is normalized to one, B is charging 5% more than A, C is charging 10% more than A, and D and E are charging roughly 30% more than A for the same product. And those price levels seem to line up with the frequency of the price effects.

But it's not just that D and E change their prices much less often and only do it on Sunday mornings occasionally. They're also just setting higher prices. And then B and C, which are changing prices more often than D and E but not as often as A, they're setting intermediate prices on average. And so motivated by these findings, what the paper then has is it's got this theory section looking at saying this question is what prices would we expect if firms had different pricing technologies and some were faster than the others?

And both A is best reacting to D and D is also optimally pricing, knowing that we're going to set a price and then A is going to undercut us and then we're going to be stuck for a month being undercut by A. How should we price? And what they argue is that if you think about that, even just think about the static game. We're going to play this game just once. I'm going to set my price and then you're going to get to undercut me and I'm going to be stuck because I have an inferior updating technology. How should you price?

And what they argue is that if you think about a differentiated product like Hotelling model or logit demand model, what you're going to do is the Stackelberg leader here, the person who moves first is disadvantaged, because you set your price first and then the other guy undercuts you. But what are you going to do? You're going to price higher than the Nash equilibrium price knowing you're going to be undercut, and the equilibrium is going to be where the firm that moves slow moves first, prices well above the Nash equilibrium, trying to induce the firm, undercutting it to also undercut to an above Nash price.

And we end up with two above Nash prices where the initial early mover is disadvantaged, but it's disadvantaged, but it sets a high price to try to induce the other firm to go high underneath it. And knowing that you're going to be undercut causes you to price higher than you would in a simultaneous move pricing game. And they argue that this asymmetric speed of adjustment could be something that leads to super competitive prices where, in their model, even A is pricing above the static Nash point. But it's like D and E are pricing well above static Nash to try to pull A up knowing that they're going to be undercut, but at least they're being undercut and getting a moderate market share at a higher price.

And again, like in pure Bertrand, this would not work well. But if you think about this as like CVS competing with Amazon, most people buying from CVS are just going to cvs.com. There's got to be some cross elasticity between cvs.com and Amazon. But a lot of people just buying Zyrtec at CVS just go to cvs.com, they buy their Zyrtec. And so if you have these finite price elasticities, they argue this pattern is something you might expect to see of slow price adjustment being a factor that can lead to higher price levels.

Any questions? So that is it for today. So, yeah. Next week's topic is entry. Again, entry theory, Monday, empirics, Wednesday. And Monday's lecture, really, it's a lot of textbook material. I will be covering the Nikhil Vellodi paper towards the end, but it'll mostly be stuff from Tirole's book.