[SQUEAKING] [RUSTLING] [CLICKING]

GLENN ELLISON: So anyway, today's lecture, I was doing-- I did static competition last week. Today, my plan is I'm going to start with empirics of static competition. I'm going to fill in, in the middle, some extra theory that I didn't have time to do last Wednesday. And then the paper at the end, I'm going to do very, very quickly because I just was reading it over, making up the notes, and I realized that Tobias is going to be spending an awful lot of time next Monday explaining what the model is about and the estimation of it.

And I figured that it'd be much better to wait until after he's done that and you can just look at my slides on it. And I'll just stick with basic comments. So anyway, I'm going to first cover one of my favorite papers in the course. So this is Bresnahan's paper, "Competition and Collusion in the American Automobile Industry: The 1955 Price War." I should say, in some sense, it's one of my favorite papers. I took the star off it this year because in some ways, I think as the paper gets more and more dated, you can appreciate it more by hearing about what it does than by actually reading it.

There just many things that you read and you're like, wow, that sounds like a 40-year-old way to do something. And so I'll just leave out the 40-year-old parts that were a bit awkward and probably had some trouble publishing it to begin with, and I'll give you the main ideas. So this is a paper on the IO - Economic History boundary.

The applied topic is what caused the 1955 surge in US auto sales, and it has this interesting hypothesis, which is that the 1955 surge in auto sales was due to a breakdown in collusion that, at that time, there were very few foreign cars sold in the United States. There were largely three American automakers. And his hypothesis is that the three US automakers were colluding to set monopoly prices, and then in 1955, the collusion broke down, prices dropped, many more cars were sold.

Anyway, and then the firms got it figured out again. In 1956, prices went back up. The paper, besides having an interesting historical fact, it has been really influential as a model for future work. In some sense, I see this as the first structural empirical paper in IO, and I think that what you can think of is that really many future papers have picked up on Bresnahan's methodology and adopted this as a different way of doing things. So I will explain what it was that Bresnahan was trying to do in this paper.

But first on the applied topic, you know, here's-- he starts with a simple table of car sales by year in the United States, and you get this striking finding, which is in 1954, US automakers produce 5.5 million cars. In 1955, they produce 7.9 million cars. In 1956, they produced 5.8 million cars. Cars are a big part of the economy, and the cars being produced going up by 40% in one year is a really striking fact. If you look at it, sales in dollars, it's still a big jump. The jump in dollars of car sales is not as big as the jump in unit car sales.

In part, that's because prices were lower, and in part, that's because people did buy smaller cars. And then if you look at what was going on in the US economy, it was a boom time in the US economy. You can see non-auto durable goods spending went from 14.5 to 16 billion and then continued increasing to 17 billion the next year. So it was a good year for the US economy, but not a 40% larger increase in car sales boom in the US economy. There really are not 40% boom in car sales general macro shocks.

You can find explanations in the popular press of what caused the 1955 thing as a cultural phenomenon. Here's a quote from Car Talk describing this, where they talk about how rock music and drive-in restaurants were becoming a thing and there were sleek new designs for cars and whatever. Apparently back in the 1950s, Paul Samuelson taught both [14.]121-- taught both in the micro sequence, [14.]121 through [14.]124, and in the econometric sequence, [14.]381, [14.]382. I think the numbers were still the same 50 years or 70 years ago, that is now.

But Paul, apparently-- this was pre-14.192, but there was a required second year econometrics paper. And Paul Samuelson apparently was teaching the second year econometrics course, would say that he would flunk any paper that claimed it was providing an explanation for 1955 sales, where by "explanation" he meant you're regressing sales and a bunch of right-hand side variables like the price of steel and whatever. OK, so what were cars in 1955?

So I've got some illustrations here. So 1955 Mercury Montclair, 1955 Chevy Bel Air, 1955 Ford Thunderbird. You can see the '55 Thunderbird is a very cute car. The other cars just look pretty generic to me. For contrast, here's the 1955 Chevy Bel Air versus the 1954 Chevy Bel Air, and you can see that they seem to have taken away a little bit of this wheel thing that's stuck out here and they smoothed out the grille.

Instead of these big things, they went to small things. You know, it fits with the sleeker car theory, but it doesn't look all that different to me. Perhaps someday-- perhaps today, someone could write a machine learning thing judging the sleekness of cars and seeing how much the sleekness of cars affects increased sales or something like that. I don't know.

OK, anyway, so his hypothesis is that this was caused by a breakdown in collusion among the big three. In general, if you're writing a-- there are several types of empirical papers one can write. If you're writing a, here's a fact. I have an explanation for that fact. Generally, if the fact is well-known and you're proposing an explanation for it, there are a thousand explanations for any fact you can come up with.

So generally, the way you try to make the case that your explanation is correct and convincing is to think of a lot of other predictions that go along with your explanation that people haven't talked about yet, and then show that here's a bunch of other implications of my theory. Now I'm going to look at the data in more detail and show these other implications of my theory are true as well. And that's what Bresnahan is trying to do.

And his idea is that motivated by then recent literature on differentiated product competition was that if there was a shift from collusion to competition, it ought to have a different effect on different car models. In particular, car models that were close in product space to other car models would see a much bigger drop in price than car models that were very distinctive. So on my previous slide, you can see the Ford Thunderbird. It was an original car. It looks very different from those other cars.

Maybe if you're colluding or in competition, it doesn't matter so much for the price of the Ford Thunderbird. You still have a lot of market power. But at least to my mind, if you're looking at the Mercury Montclair versus the Chevy Bel Air, I would have to spend some time learning to recognize the difference between those two if you were showing me pictures of them. And so you might think if there was Bertrand competition between Mercury and Chevy, the price of the Bel Air and the Montclair, they could be very tough competition and drive prices down towards cost.

So just, Bresnahan has these pictures illustrating this. He supposes that, you know, suppose that this was just a pure one dimensional vertical competition model like I described last time. So we put quality on the x-axis. This solid line is supposed to represent cost. And then these are the prices of the different car models where General Motors and Ford are the two firms producing them. It is true that generally, General Motors tend to produce the bigger cars and Ford tended to produce the smaller cars.

So GM, with its Cadillac division, had a lot of the upper end of the quality spectrum more to itself. And what he says is that in this vertical differentiation model, if you have models that are relatively far apart, like models four and five, if you view these as very different in quality space, then if these are the collusive prices and they start competing with each other, it may be that prices for four and five don't drop very much because they're still far enough apart that the marginal consumers are small relative to total demand so you don't do much to try to attract more marginal consumers.

But on the other hand, if you look at models two and three that are very similar-- so this could be my Bel Air and Montclair. If they're colluding, they may charge fairly high prices. And then if they start competing with each other, prices may get driven very low because they're very close in product space. And so if you had that as an idea and you thought about how you could test that, Bresnahan could have taken this sort of, what I'd call the classic reduced form approach in IO, which is you make predictions, you test the predictions.

And so what you could have looked for is, you know, are price drops from 1954 to 1955 larger for car models that are closer in product space? And that would be an auxiliary prediction of a breakdown in collusion. But, you know, Bresnahan thought that this is really not all you could do with that theory. One thing is that the theory is more complicated than that. Like if you remember the vertical competition model, every firm does have a first order condition that's pinned down to the firm above and to the firm below, and those two first order conditions are going to pin down the prices.

But those things are going to involve the prices of the next lower quality model and the model below that and the model below that. And so it is a system of n equations and n unknowns, and any change in the competition-- like the change in any one place is going to have ripple effects throughout the distribution and will change the prices of the higher price models even if they're not in very close competition. And then the other thing that Bresnahan said is that our theories are quantitative theories, not just qualitative theories.

They don't just predict that prices will go down. They really are predicting how much the prices would go down, and how much the price would go down depend on things that we, in theory, could measure. And so really, to give a more complete test of the theory, we should really see whether the theory can quantitatively account for everything instead of just qualitatively account for everything. So what did Bresnahan do? And this is where the paper starts to look awfully primitive these days.

He basically wrote down the standard vertical differentiation model that I described in lecture. So suppose there are n models with qualities v1 through vn, and suppose that the quality of a car is simply a function of observable characteristics. What function of observable characteristics? Basically it's just x times square root of x times beta, where x would contain things like the horsepower of the car, whether it had a V8 engine, the size in cubic inches, the weight of the car, whether it had air conditioning, things like that.

He then assumes that there's, again, just like in the vertical differentiation model, there's a mass of consumers with types theta distributed uniform on some interval zero to vmax, and there's a density delta on that interval. And then he says that the utility of a type theta consumer from buying model i is going to be theta vi minus pi. So pure vertical model. And then he has the cost of producing q units at quality v is some constant or some function of v that's going to end up being irrelevant, and then q mu e to the v power.

So costs, again-- v somehow is like the square root of the observable variables and then costs are exponentially increasing in v as you try to raise quality, and then there is-- it's a constant marginal cost. So mu times e to the v is the marginal cost of producing a single car. But then what Bresnahan's thought was, is that this is a model that I can fully specify. And if you fully specify the primitive-- if you fully specify the primitives, which is the production functions of the firms, the preferences of the consumers, and the model of firm behavior, then you can fully predict what's going to happen.

So basically once you fix the parameters, beta mu and the production function and delta in vmax, in the consumer preferences, you can fully predict what's going to happen under any model of behavior. And so in particular, collusive pricing just becomes a multi-product monopoly model that we covered. Static competition would be another natural model, what people often apply to a situation like this. And then Bresnahan also considers two other models.

One of these is like non-cooperative static competition between all products. So this is as if every Ford model is run by a different manager who's competing aggressively against every other Ford product. So every product in the market is run by a different manager who's profit maximizing solely based on that car. And then he's got a fourth one. This is one of those ones where it's based on somebody's paper from the 1970s, and I have a hard time-- it's sort of perfect competition. You have a hard time making out today what is that model and what would be the justification for it.

But he's got four different models. He notes that under any of those four models, you can work out what should happen given any set of parameters, so then you can invert that relationship and say, given any one of these models and one of these models applied to any-- for any behavioral model in any year, 1954, '55, or '56, I can figure out what would be the parameters that best explain the data. And then Bresnahan's idea was for each model of firm behavior in each year, I figure out what parameters best explain the data and then ask how well does that model explain the data with those parameters?

And then I pick out whichever model fits best in every year to say this is how the firms behaved in that year. And the way he does this is-- again, this seems archaic, but it's a set of non-nested hypothesis tests where he looks at how well the-- looks at the model, fits the model to the data, and then, in a likelihood ratio sense, says, do the predictions of the other model help explain the errors? Do the predictions of model two help explain the errors in model one?

And if predictions of model two help explain the errors of model one, then model one is rejected by model two and model one is not the correct model. And so what he looks for is, is there one of these four models-- in each year, year by year, is there one of these four models that's not rejected by the other model? And then the other thing that Bresnahan really also pioneered, and this was a paper in his thesis, was the idea of a counterfactual.

His thought was that if we do this and we figure out what is the correct model in each year, then I can do a counterfactual and predict what would have happened under any other model or any other government policy. I could tell you what if the government had had a 5% sales tax on cars, what would happen if the firms had colluded instead of competed with each other, what would have happened in-- any other thing change you can think of to the industry, you can predict what would have happened under his model, and it gives you the counterfactual prediction.

OK, so what are the main applied results of the paper? And I always think of it like maybe someday, someone will rewrite this paper. Obviously this is primitive data, like minimal data from the 1950s. Maybe someone will rewrite this paper and say, from a more modern perspective, what do we think actually happened? But what does this analysis say? First one says that the collusive model fits the data best in 1954 and 1956, and the static competition model fits the data best in 1955.

So that's his evidence for his main hypothesis that there was this breakdown in collusion and the collusive model fits the patterns best in '54, '56, the competitive model in '55. Further supporting that is his model has a bunch of parameters, the mu and the delta and the vmax, and especially the betas that talk about-- which relate to how valuable is horsepower, how valuable is air conditioning, things like that.

What he notes is that if you estimate the collusive model in '54 and '56 and then estimate the static competition model in '55, the parameters like the betas turn out to be fairly similar between '54, '56, and '55. And that's further supporting evidence that consumer preferences towards horsepower should probably be fairly stable over time. And so what he notes is that if you do this model in these two years and the other model in '55, then the preference parameters are fairly similar across the years.

Whereas if you try to fix any one model and use any one model to explain what happens, the parameters need to change a lot from '54 to '55 and from '55 to '56. In some sense, you need-- prices did go down despite the fact that demand is going way up, and so you need something to say that suddenly cars got a lot cheaper-- cars got a lot cheaper to produce and people valued cars more than they used to value cars during that one year.

Whereas if you use this change in behavior, then the parameter estimates are fairly stable across time. And so that's supporting evidence. Then what's in his counterfactual? The counterfactual he does is suppose firms had colluded in 1955 instead of engaging in differentiated product competition. And what he finds is that prices would have been 6% higher and demand would have been 11% lower. So this was a 40% jump in auto sales.

His model is not fully explaining it. His model is saying, I'm explaining a little more than 25% of the boom in auto sales as coming from the change in collusion. So in some sense, he's already letting you know that his model-- at least the model as he's estimated-- isn't fully correct, because it's not fully accounting for the jump in sales. Now, there was a 10% jump in non-auto durable goods sales. You would have said just the background change in US economy income or change in income maybe could have explained 10% also.

So he's saying his model is explaining part of what's going on. But it is a very nice use of the counterfactual to give us a sense of what would have happened in another world. You see every IO paper now has counterfactuals, which I think started here. As I said, if you read the paper, there are all kinds of things that are going to seem odd to you, like the adoption of the pure vertical model. Clearly, I don't think the Mercury Montclair and the Chevy Bel Air are so similar that there would be purely Bertrand competition between them.

If you thought about doing that in today's car market, there are pickup trucks. There are convertibles. There are sports sedans. There are electric vehicles. If you try to put them on a one dimensional quality space, some of them, some pickup truck, is going to nearly coincide with a Honda Accord. Clearly, you need to have some kind of horizontal differentiation as well as vertical differentiation.

Anyway, a number of aspects when you see how people do modern empirical work that look very different. But I think in some ways, the basic ideas still seem very spot on. Particularly, the one thing that I think he does better than many current models is he does take very seriously this idea of, I want this empirical-- this structural work only works if I choose the right framework, and he's going to come up with four different models and put a lot of effort into testing the four different structural models against each other.

Whereas now very often it's like it's hard to estimate even any one model, so you see a lot of papers that have, here's my one model, and I can estimate that and I can talk about how well it fits by putting up a graph and showing some x's and y's and show the y's are close to the y-hats. But I think Bresnahan's doing a very nice job on that dimension. So any questions on Bresnahan? Yeah?

AUDIENCE: Why would collusion break down for just one year and then all of a sudden, the next year it's back at it again?

GLENN ELLISON: Yeah. You know, a little hard to say. I will, in a couple weeks, be discussing models of repeated game collusion. And in models of repeated game collusion, you do get this phenomenon of the way you support collusion is via price wars to punish somebody who cheats so that as if you're trying to collude, you think that someone is cheating on the collusive agreements and their dealers are charging-- you've all agreed that we're going to charge $7,000 for our Bel Air and Montclair, and you think the other guy's dealers are actually giving consumers cuts and offering it for $6,500 instead of $7,000 even though the list price is $7,000.

You might decide we're going to have to punish you for doing that and we're going to abandon colluding this year, and we're going to all learn how bad it is and then go back to cooperating the next year. You know, so you could tell some story like that. There's no discussion in the paper of why collusion would have brought down-- there's no narrative evidence at all of the executives talking to each other and having anything going wrong.

But it is true that actually-- we're going to study several other papers on cartels. Almost all the cartels we study are finite lived. They go on, they're colluding with each other for some number of years, and then it breaks down at some point. In part because also it's illegal, but it's-- so people are afraid of getting caught or something goes wrong and they can't discuss it with each other and get things back. Anybody else?

OK. So second thing I was going to do, then, I was going to jump ahead before I get to Miller and Weinberg and just do a half hour, 40 minutes on oligopoly price discrimination. And I think in Tirole's textbook and in many other places, price discrimination is typically taught as a monopoly topic. And I think that was because historically you had monopoly and perfect competition as the two models. With perfect competition, price equals marginal cost, you can't price discriminate.

But with imperfect competition, price isn't equal to marginal cost. There's no reason why oligopolies can't price discriminate just like monopolies price discriminate. And particularly, there was an early paper in this, Borenstein and Rose, 1994, *JPE,* look at price dispersion in airline markets. And what they find is that if you look at airline city pair markets, like Boston to New York, New York to Phoenix, whatever, the markets that have more competition tend to have more price dispersion on the same plane.

So if you look at a flight from Boston to New York, there are a lot of people on that flight. They all pay different prices. And the Borenstein-Rose finding is that the heterogeneity in prices on a flight is actually larger in the more competitive markets, not smaller in the more competitive markets. So that's an indication that it's not the case that competition that we often see makes it impossible to price discriminate.

OK, so how would we model price discrimination? Well, let me first start with third degree price discrimination. So you have two firms competing in two separate markets, and suppose there's no cross market arbitrage. So this is two firms selling cell phone service in Boston to the same two firms selling-- or two firms selling cable TV service in Boston and selling cable TV service in Philadelphia. You can't get buy cable TV in Boston and then use it in Philadelphia or vice versa. You have to buy the service at your service address.

And so you could have these same two firms competing and they're competing in these two different markets. So suppose that each of the markets has Hotelling preferences and the type theta consumers in market m get utility v minus p price of firm one in m minus tm theta if they buy from firm one and v minus p 2m minus tm1 minus theta if they buy it from firm two. So this is just competition on two lines.

This is market one and it has some parameter t1, and this is market two and it has some parameter t2. I think I made t2 the high guys. OK, so you just have-- you're competing on this line, you're competing on this line. These markets are completely separate.

Price discrimination is just like third degree in monopoly. It's just two totally separate markets. So if we have competition here, what's going to happen is you're going to get p star equals c plus t1. You have competition here, you're going to get p2 star is c plus t2. We just compete on a line here. We get Hotelling here, we compete on this line, we get Hotelling here. The prices are different across markets. We're discriminating against consumers in this market.

The high types end up paying more. What are high types? A high type is somebody who lives in a market where there's more heterogeneity in taste. So if there's more heterogeneity in taste in market two than in market one, then people in market two end up paying more than people in market one. OK.

The one policy question you get is, well, what would happen in a market like this if we banned people from-- what would happen if we ban people from price discriminating and say, you have to charge the same price in both markets? It turns out, in this model you just get an average markup. So if instead of being allowed to price separately, you force them to charge the same price, what you get is price equals c plus 2t1t2 over t1 plus t2. This is what mathematicians refer to as the harmonic mean of t1 and t2.

And basic math fact is that if t1 is less than t2, then t1 is less than 2t1t2 over t1 plus t2, which is less than or equal to t1 plus t2 over 2 less than or equal to t2. This is called the arithmetic harmonic mean inequality. You can think of this as this is the reciprocal of the average of the reciprocals of t1 and t2. So it always lies between t1 and t2, and it is less than the arithmetic mean of t1 and t2. But just the way the math works, this is the price that you end up charging if you ban price discrimination.

So what does banning price discrimination do? Typically, it makes the high types better off. Typically it makes the low types worse off. It does eliminate a misallocation of markets. In many models, if you have a high price in one market and a low price in another market, there are people in the high market who would like the good who aren't getting it and there are people in the low market who don't like the good as much who are getting it. That's a misallocation like there is in price discrimination models.

So banning price discrimination eliminates that inefficiency. That inefficiency isn't present in Hotelling because in Hotelling everybody buys, at least if he is big enough. And then welfare can go either way. Here, again, there are no welfare changes. It's just a transfer because everybody buys-- all efficient transactions take place. First best is achieved in both models, so it's just a pure transfer. But in general, banning price discrimination can increase total social welfare or it can decrease total social welfare.

OK. There's a set of papers in the 1980s noting that this is the typical effect of banning price discrimination in oligopoly markets, but the typical effects don't have to hold. So there's a paper by Thisse and Vives, *AER* '88, that makes the point that allowing price discrimination can make all consumers better off. So price discrimination can make everybody better off because of this effect that Thisse and Vives highlight.

When you let firms price discriminate, it lets firms aggressively target a rival's natural consumers, and that can intensify price competition. So the simplest example that Thisse and Vives gave was consider again the Hotelling model. And I'm going to give the Hotelling model a different interpretation. So we have these two firms and they're competing, and the consumers have some types theta. And I'm going to assume that the firms can observe theta.

So maybe these are two firms selling lumber to home builders and they know where the home builder's jobsite is, and so they know how long it's going to take the home builder to truck the lumber from their warehouse to his factory site so they know his theta. The theta is determining willingness to pay through some means that's known to the sellers. And so here, discrimination would be firms set theta dependent prices instead of theta independent prices because they observe theta.

OK. Well, first, what happens in this model if firms are-- if they're not allowed to price discriminate, then you just get p star equals c plus t. Because if I know the locations, but I'm not allowed to charge different prices, then I just charge c plus t to everybody. What's going to happen if I now start discriminating? Well, if I'm firm one or if-- let's say I'm firm two and I'm trying to sell to a consumer over here. I know that that consumer prefers firm one, prefers firm one's product to my product.

To get that consumer, I'm going to have to undercut firm one. In pricing to this consumer, there's no product differentiation left. We know this consumer has some value v1 for firm one's product and some value v2 for firm two's product, and we're just pricing to that consumer like in a Bertrand competition model. So what's going to happen is Bertrand competition is just going to force prices down until p2 star equals c and p1 star is going to be-- so p2 star of theta is c and p1 star of theta is going to be firm one's advantage, which is negative theta minus 1 minus theta.

Minus 1 minus theta. So that's 1 minus c plus. So c plus 1 minus 2 theta. So what firm one is going to do is firm one is going to price at its advantage relative to firm two, and the consumer is going to buy from firm one. And so what you get if you-- over here I've tried to graph what is equilibrium price paid as a function of theta, and all these consumers over here buy from firm one.

As you come to the middle of the line, at the middle of the line they're in pure Bertrand competition. The consumer has the same preferences for both, so they both price at c. And then you go back up here to c plus t at the end. So no discrimination. Everyone pays c plus t. With discrimination, the worst case is you pay c plus t if someone knows you're fully captive and like their product by t better, and then otherwise you pay c in the middle.

OK. So all the consumers are better off, and they're better off because the rivals are cutting prices to target consumers who don't want to buy from them, and that targeting the rival's consumers is an intense-- allowing them to do that intensifies the competition, and that intensified competition cuts prices for everybody. So that's the thought that price discrimination can be a universal good for consumers. OK.

A second insight, though, came, Ken Corts, a decade later, was that that's true. But there's a second effect. And price discrimination can actually make prices higher for all consumers. And it's the thought that price discrimination makes your pricing more flexible. One thing that lets you do is target your rival's consumers, and that's good for competition. The other thing it does is let you exploit people who are captive to you.

And so Corts looked at this specification where if you think about it-- so over here, what I've tried to graph is-- what that's a graph of is what is v, the value that you get for buying from firm theta. And what he assumes is that you get value two from buying from a consumer-- so these are value they get buying from firm one. Your value goes like this. The farther your preferred type theta is from firm one's product, the lower is the utility you get from it.

And then once the product is too far from what you want, your utility just drops to zero. So I have this decreasing distance from firm one. I like it less and less. And then once it gets to some point, I'm like, I can't use that. And then your utility drops to zero. So that's my utility buying from firm one, and my utility buying from firm two does the opposite. It's zero here and then it goes up like that. So these high theta consumers really like and get utility two from buying firm one's product. The guy in the middle gets utility one from both, and vice versa.

So in some sense, I've eliminated that targeting your rival's competitors effect by saying that anyone who gets utility from one firm gets zero utility from the other firm. So letting them target your rival's consumers doesn't actually create any competition. So what would happen in this model? In this model, if I ban competition-- if I ban price discrimination. Sorry. If I ban price discrimination, firm one is essentially pricing as if it had a linear demand curve.

It's a linear demand curve that cuts off, but if you remember, if you have a linear demand curve that looks like this, even if it continues, you still would price at one. If you ban price discrimination, firm one is going to set p star equals 1. Firm two is going to set p star equals 1, because they're both just pricing to the linear demand curve 2 minus p. Whereas if you allow them to discriminate, I do first degree price discrimination. This guy pays 2, this guy pays 1.9, this guy pays 1.8, and so on until this guy pays 1.

So basically, I just raised prices to everyone because letting me price discriminate lets me exploit all-- lets me extract all of their surplus when I don't have competition, and here I've designed the oligopoly so there really isn't any competition. So I think Corts's basic insight was that there are these two effects and we shouldn't necessarily expect either of those two effects to dominate. Price discrimination can be good because it intensifies competition.

It can be bad because it has more exploitation, and it's just going to matter on the distribution of relative preferences for firm one versus firm two's products, whether the exploitation effect dominates and makes prices go up or the competition effect dominates and makes prices go down. As I said, this Hotelling model, it's like it's an intermediate effect. It's like you charge intermediate prices and you lose the ability to compete and discriminate in a way that makes you just charge something in between what you would have charged the high and low types. Any questions on that argument? Yeah?

AUDIENCE: He's just kind of arguing that if it's not actually an oligopoly, it's monopoly.

GLENN ELLISON: Yeah, that's right. So he made it a monopoly, but then obviously you could have it be an oligopoly like this and you would get the same results, Like if it was an oligopoly where it really was an oligopoly but most of them got low utility from one and high from the other, then you would still get the same effect there. But he's arguing that it's basically it's adjacent to the monopoly model. In the monopoly model it goes this way, so oligopoly model is near that one, it's going to go the same way. Yeah.

OK. So that was oligopoly third degree price discrimination. What about second degree price discrimination? And I do think that oligopoly second degree price discrimination is a really important topic because I think many, many products people buy, there is second degree discrimination there. You're buying phones. There is the iPhone 13 and the iPhone 14 and the 128 megabyte and the 512 megabyte version. You're buying cable TV and it's cable TV package one, package two, package three, package four, all with increasing features.

You have competing firms selling TVs. Everyone's got the 42-inch TV and the 49-inch TV and the 55-inch TV and the 62-inch TV, and each of those TVs has the 4k or doesn't have the 4k. But you generally see firms in oligopoly industries are selling multiple products, competing with multiple products of each other. So if you're talking about oligopoly second degree discrimination or the cars in Bresnahan's example, like it's Ford is producing a line of cars, GM is producing a line of cars. Those cars are competing against other models, and there's vertical constraints within the firm and there's also competition across the firms.

So clearly you want to have vertical differentiation and you want to have horizontal differentiation. And then the substitution patterns, there's going to be cross firm substitution going on and there's going to be within firm substitution going on. And it just turns out it's hard to write down a nice tractable model that has the two dimensions of consumer tastes and has both substitution patterns mattering and the binding IC constraints and the binding cross firm constraints.

So what the literatures tend to do is do a few different special cases. One thing it's done is to assume independent horizontal and vertical preferences. There's this observation that if you make horizontal and vertical preferences independent, you get a surprisingly tractable model that you can do even with a continuum of types. The other thing that people do is have discrete vertical types, and when you have discrete vertical types, you can sometimes make the IC constraints, the within firm IC constraints non-binding and eliminate the number of constraints you have to solve for.

And then a third approach you sometimes see people do is actually just eliminate the vertical differentiation entirely and talk about second degree discrimination in a model that doesn't have vertical differentiation. So let me go-- I'll start with just an overview of what happens in this independent model. So Stole has a chapter of the handbook in IO that surveys this literature that's pretty big and seemingly general. You have these firms j. They can produce a range of qualities s at some constant marginal cost c of s. So typically they actually have a continuum of s's they could produce if they wanted to.

Consumer i's utility from buying a good of quality s from firm j is theta i of s. So this is the vertical preference. High types are willing to pay more for quality. Minus the price plus the epsilon ij. So this is the logit-like error that we've had in our models. And then models assume that the theta i's and the epsilon ij's are all independent, and theta i is distributed f theta, epsilon ij is distributed f epsilon, and they're pretty general in terms of what distribution of theta is and what distribution of epsilons you have.

So let me define s star of theta to be the efficient quality for a consumer of type-- that should say type theta to buy. So you have a consumer. That should be a theta instead of a v. So if you have a type theta consumer, they ought to be buying quality that maximizes theta times s minus c of s. So this is if the goods were sold at cost, what would they buy? That's the efficient thing to sell to them.

So just an example of a model fitting this. Consider a Hotelling-like model, and suppose that epsilon i1 is distributed uniform on zero to t and epsilon i2 is just the negative of epsilon one. A Nash equilibrium of that model is just pjs is c of s plus t for all j and s. So if you imagine, it's as if we've got this-- I've got this continuum of Hotelling lines. So you have these different quality levels, s1, s2, s3, s4, s5.

On each line, let's imagine I charge p star equals c of s. So quality of-- yeah. Price for quality s is c of s plus t. So imagine you had that pricing. That every quality I sell at the cost plus t. Then what are consumers going to do? The markups are the same for every quality level. So when I'm deciding which quality level to buy, I buy the efficient quality level. If I'm a type theta consumer, I buy s star of theta. So all the consumers for whom s3 is the optimal quality are like, OK, I'm either going to buy quality s3 from firm one or I'm going to buy quality three from s2, who do I buy from?

Well, again, it depends on my-- now I'm calling this an epsilon i1 and an epsilon i2. But if my taste is I really like firm one's product, I buy from firm one. If my epsilon is I really buy firm two's product, I buy from firm two. And so it's just the-- as far as if the firm cuts its price, it's just like the competition on a line model.

And if it cuts its price, it will attract more consumers from the other side. It would have this increase in demand one half plus p2 minus p1 over 2t, but they're just basically competing on a line against the other firm for this quality level. And so the same Hotelling first order condition applies and they charge price c plus t on each line. And so if the firms are doing this and they're playing Hotelling on each line, they're pricing at c plus t on each line, all consumers do buy the best quality for them.

And so the result here is that when you have competitive secondary discrimination, secondary discrimination goes away. The firms are not using different quality levels to discriminate among consumers. They just are charging-- it's not like the iPhone with 512 megabytes cost $300 more than the iPhone with 128 megabytes. They give them all at the same markup and all consumers buy the efficient quality. So all of that quality distortions in secondary discrimination goes away.

And I gave this example in the Hotelling model, but Stole shows that this is really very general. So there's this very general thing of for f distributions of epsilon, satisfying some conditions, industries of epsilon satisfying some conditions, as long as the theta and epsilon are independent, this always happens. You just get competition on many lines and constant markups across goods.

And I think the thing that I want to point out most about this model is that it seems very general, but it's got this hidden, special, and incorrect assumption in it, which is that the thetas and the epsilons are going to be independent. And a good example to think about this is think about you and a CEO of some company flying to Europe, and you're told that there are two flights available to go to Paris. There's a Delta flight or an Air France flight.

You might say that some people like Air France better than Delta. Some people like Delta better than Air France. It may be that the preferences that students have on Delta versus Air France and the preferences that the CEOs have on Delta versus Air France are uncorrelated. That yeah, they're both-- like, you can model them as uniform on zero to t for their disutility for Air France relative to their ideal airline disutility for Delta. But what's uncorrelated is reasonable. Independent is not reasonable, because how much are you willing to pay?

Those of you who really like Air France, how much extra are you willing to pay for to go on Air France? Maybe it's $20. Maybe it's $10. I don't know. The CEOs, they may pay $5,000 more to go on Air France relative to Delta if they like it, because for them, the $5,000 is rounding error. And so I think what's reasonable there is the uncorrelated preferences that the businessmen and the students may have the same-- well, same uncorrelated preferences for Delta versus Air France, but they wouldn't have independent preferences.

What it really might look like is you'd have the student line-- you have the student line where t is very small, and then you have the businessman line where t is very, very big. So they have a higher vertical willingness to pay for the good, but then, in addition to the higher vertical willingness to pay, they just have much more intense horizontal preferences as measured in dollars. That's not covered by the independence case. Independence would say there should be one line up here like this where, again, they'll all shift from Air France to Delta for 20 bucks, which is just not going to be true. And it's the independence, in some sense, that results in the constant markups.

So a second way people model this is just let's go with this discrete type space. So there's just two quantities, sl and sh. sl is the coach class seat where your knees are pressed into the seat in front of you and then sh is the lie flat bed with the nice TV and the champagne and glasses when you walk on the plane and everything else. And so what you can think of here is model this is I can model that as the vertical differentiation model where consumers have a discrete vertical type tl or th, and then they have horizontal types epsilon i1 and epsilon i2, just like in the Hotelling model.

And let's instead then specify the utility. Instead of being theta s plus epsilon minus p plus epsilon, we'll have the theta multiply s minus epsilon. So that if you have-- everyone has horizontal preferences at the same intensity in utils. But when you convert the utils to dollars by multiplying by this theta, the preferences in dollar terms are much more intense for the businessmen than they are for the students. So the theta, the vertical type, multiplies both the quality and the relative preference of Delta, Air France. OK.

So in this model, it's reasonable to assume that the thetas and the epsilons are now independent. So you have this willingness to pay for quality, whether it be quality of the nice seat or quality of the Air France food. And then you have an idiosyncratic taste for the Air France food versus the Delta food. OK. Anyway, so I'm going to assume for this numerical example that the thetas and the epsilons are independent and half the people are high types.

So what happens here? Well, you know, simple fact. I've got conditions under which this is true, but under some set of conditions, what you get is if this is a line that has length tl and this is a line that has length th, you get p star equals c plus tl here and you get p star equals c plus th here. All right. So on the big line, you price-- with the big amount of differentiation on the little line, you price in the little amount of differentiation, and all the low types by the low quality product, all the high types by the high quality product.

What has to be true for that? It has to be that the th times delta s is the willingness to pay of high types to get the high quality product. tl delta s is the willingness to pay of the low types to get the high quality product. And it just has to be that the markups that result from different degrees of differentiation are exactly between the willingness to pay of the high types and the low types. So here, just in this case, under those assumptions, the IC constraints are all non-binding.

There are no consumers who are indifferent between buying first class ticket or a coach class ticket in this model I've written down. Students are just in competition on a line to buy the low class tickets. These are contrasted by the high cost tickets. It requires a coincidence. The coincidence is that th minus tl is in this range that makes both of the IC constraints non-binding and makes the businessman buy the business class tickets, the students buy the coach class tickets.

I don't know. This is a perfectly reasonable model of secondary oligopoly price discrimination. Again, it does say that competition determines markups. So it is a very different mathematical structure from what you get in the-- it's a different mathematical structure from what you get in the second degree monopoly price discrimination model, where it's the IC constraints and the manipulating quality to do things. It is the case that you can still get-- you can think of you can get-- I set these sh and sl exogenous.

It could be that you can get manipulation of the qualities here where firms charge-- choose a quality level sl that prevents the businessmen from taking-- in some sense, they choose how far apart to make these lines so that this price discrimination scheme holds up. If sl were too high, I increase sl at some point, the IR constraint would start binding and the high types would jump down and start buying the low quality good instead.

And we say interesting about this model is this model is also a model of multiple equilibria. Sometimes this model has two equilibria, one where the firms discriminate and one where the firms just sell the high quality good to everybody. In some sense, there's just this complementarity. We know these things have upward sloping reaction functions where you tend to-- you only go one half to one, but as your rival raises their price, you raise your price. That same thing means if your rival is spreading its prices for the high and low quality good apart, you want to spread your prices apart.

So you can actually have two equilibria, one where the low quality good doesn't exist in the firm. They just compete for everybody selling the high quality good, and then the other where this discrimination occurs. But the discriminatory equilibrium is better for the firms than the one where they don't discriminate. Any questions on that model?

OK. So then the final theory topic I want to get to is a little bit on add-on pricing. So at first let me go back to Lal and Matutes *Journal of Business,* 1994. Lal and Matutes were thinking of this question of, why is it that turkeys are on sale on Thanksgiving? Wouldn't you think that the one day of the year when turkey demand is a hundred times higher than it is any other day of the year, turkey prices would be high instead of low?

And let's think about the following model. And think of this as a model of people competing. This is competing supermarkets, and the supermarkets are selling two different products, v1 and v2, and consumers actually all want to buy both products. They want to buy the turkey and they want to buy the stuffing or the cranberry sauce or whatever. And so what my utility is, I get v1 plus v2 minus p11 minus p12 minus p theta.

If I am type theta and I drive to store one and I buy both products, they are at these prices. Or I can drive to firm two and again get v1 minus p21 plus v2 minus p22 minus t times 1 minus theta if I drive and buy the products from firm two. And then I can, if I want, split my shopping and get v1 minus p11 by buying the turkey at store one and then get v2 minus p22 from buying the cranberry sauce at firm two. But then I have incurred both driving costs.

And so Lal and Matutes consider this a multi-stage game where what firms do is imagine there were per product advertising costs. You don't want to advertise the price of every product in your store. That's clearly true with supermarkets. You've got 10,000 products there. You can't advertise all of them. So in the extreme case, suppose you choose and advertise a price of one good, and then after choosing the price of your one good, you choose an unadvertised price for the other good which is kept secret until consumers arrive at the store.

Consumers drive to one firm, incurring costs either t theta or t1 minus theta. When they get there, they learn the price for the unadvertised good. And then consumers, after learning the price of the unadvertised good, they either buy one or both products there or they decide to drive to the other firm and buy it at firm two instead.

Proposition first. In the above model, equilibrium prices satisfy p1 star plus p2 star is c1 plus c2 plus t. Again, this is just the Hotelling competition on a line model. In equilibrium, the price of the bundle is going to be c plus t. But the individual prices are not like c plus t over 2 for each. The individual prices are p2 star equals v2 and pj1 star is c1 plus c2 plus t minus v2. So basically cranberry sauce is always incredibly expensive on Thanksgiving.

You know, you have this advertised product and the unadvertised product. You show up at the store to buy the unadvertised product and you find that, wow, that is exactly my full willingness to pay. I'm basically indifferent to walking out versus spending $11 buying cranberries. But in equilibrium, both firms are charging $11 for the cranberries. So there's the advertised price for the turkey.

The advertised price is that turkey is a great deal and it just is a great deal reflecting how much they're going to rip you off for everything else you buy when you get into the store to buy the turkey, that's a good deal. And in equilibrium, it's all neutral. In equilibrium it's as if you advertised the prices for everything, we would get exactly the same outcome. It's c plus t pricing. It's just in a distorted way where one good is really cheap and then the other goods are expensive.

And actually, there's actually a paper on Thanksgiving noting that, in fact, it is true that turkeys are cheapest on Thanksgiving than any other day of the year and everything else in the grocery store is, on average, more expensive than any other day of the year. This is what firms seem to do is, you know, turkeys are a loss leader. Actually, why are turkeys a good loss leader? Milk is often also a loss leader at convenience stores. Like a gallon of milk seems relatively cheap compared to a bag of potato chips.

Turkey and milk are good loss leaders because you don't want to have a loss leader where people come in and stock up and buy 23 of the thing that's advertised as a great deal and then not buy anything else at your store and exploit you and save it up and live off that for the next six months. Like ramen noodles would be a terrible loss leader because you'd get the student comes in, buys the case of ramen noodles, and just leaves with just that. But turkeys, you really don't want more than one.

Gallons of milk, you really don't want more than one. So something that you can give people a very good deal on, they buy exactly one, and then you rip them off for everything else, that's how you choose your loss leaders. OK. But, you know, the Lal and Matutes model says that this is how loss leaders work, but they are neutral economically. It's not like banning them would do anyone any good. They're completely neutral for the firm's profits and neutral for consumers, and they do economize on advertising costs. So in that sense, they're good.

I put Carmen's picture up there. I respect-- Carmen, at one point, was a very successful economist and then had been a professor at Berkeley and dropped out and became a novelist. And actually is a successful novelist, although she writes in Spanish so I've never read any of her books. But anyway, won several prizes for her fiction writing after she decided to have a second career.

OK. Add-on pricing. So add-on pricing thinks about basically this loss leader idea in a vertical competitive price discrimination model. So there are many cases where firms that are practicing second degree price discrimination do have non visible prices for add-ons. So for instance, you go to a hotel and it's very easy to find the price of a hotel online and compare prices of hotels.

It's very hard to compare what's the price of the beer in the mini-bar or what's the price of the sparkling water in the mini-bar, what's the price of a croissant in the hotel lobby if I'm in a hurry to get somewhere. Rental cars, it's very easy to compare the price of the rental car. It's very hard to figure out what will I pay for the car seat if I need a car seat for my infant when I get there. I don't know. I guess you guys are too young to have infants.

But it's striking to you that you go to rent a car and you can rent a $20,000 car for 19.95 a day, and then you want to rent a $37 car seat because you don't want to lug it on the plane, and it's like 13.95 a day. And you're like, you know, how can this be that they can rent me a $20,000 car for $20 and then the $37 car seat, they're charging me $13 a day or making me carry this enormous piece of plastic on the airplane.

Obviously, it's this same add-on pricing thing where it's very easy to compare the prices of the cars and not of the car seats. Therefore, the car seat is priced at v2. What is your disutility from carrying a car seat on a plane? That's what they'll charge you. And then the cars are, in some sense, the loss leader. They make money selling the car seats and the rental car insurance.

Inkjet printers-- I mentioned earlier that, again, the inkjet printer is very cheap, but they've invented this inkjet printer that has almost no ink in it when you buy it. And then it's got the starter cartridge, and then the regular cartridge of the inkjet printer is much more expensive, and that's the hidden feature when you buy the cartridge. To think about how this works in vertical competition, Lal and Matutes was purely additive competition. It wasn't a vertical model.

So imagine you have firms advertise prices for their low quality goods like the rental car without a car seat. Costs sl and sh both have cost c, as if the car seat is essentially free for them to provide. Consumers incur some cost to visit j and learn the price. So consumers have to visit the firm to learn the price of the upgrades. You have to spend some time on the website learning what it costs to get a car seat added to your rental car, and then that would let you upgrade in quality.

And then again, let's do the competition on two lines. So epsilon i1, epsilon i2, independent. Again, same thing, this setup. So this set up where what you do is imagine this set up where you're only advertising the price of the low quality good, and then consumers have to pay something after they arrive or after they've spent some time investigating your firm to learn what it cost to upgrade to the high quality good. What happens?

Well, there's an equilibrium where, like in Lal and Matutes, the price of the upgrade is the full value for the upgrade for the high types. So the price of the upgrade is what are the businessmen willing to pay for the rental car seat. And then price of the low quality good ends up being this, which I won't get into. But main applied conclusions are that profits in this game are higher than the profits in the game where pl and ph are both visible.

So in this competitive price discrimination model, add-on pricing or these loss leaders are not neutral. They raise prices. And not only do they raise prices, but they make all consumers worse off than in the game where the low quality good didn't exist. So if I banned the production of rental cars that don't come with a car seat, if you ask for it for free, all consumers would be better off. The low types would be better off and the high types would be better off.

The high types are better off because they would pay a low price without this add-on pricing scheme. The low types are better off because they're getting a better quality good instead of having to carry their car seats on the plane. I'm short on time, so why don't I skip that? But let me just give you the intuition for why that add-on pricing becomes non-neutral. Think about reducing your price to p star minus dp in the add-on model and doing the same in the model with no low quality good.

In both models, what you lose from a price cut is one half dp from charging the lower prices to the half the consumers you're serving. In both models, what you're gaining from the price cut is the number of customers gained times the per customer profit on the marginal consumer that you gain. And so because, in both models, if you tried to make-- in both models, the first order condition, this has to equal this. The one half dp is the same in both, the number of customers gained in both, therefore, the per customer profit on the marginal consumer is the same in both models.

So in both models, the per customer profit you make on the marginal consumer is the same. But profits are based on profits on the average consumer, not the marginal consumer. And if you think about what goes on in a model like this, when I cut my price, who do I get by-- I offer my rental car for $1 less than everybody else's rental car. My rental car is 38.95 a day instead of 39.95 a day. I get all the cheapskates.

And I get all the cheapskates where the low types were willing to carry their car seat on a plane and spend extra time waiting at the gate to get their car seat and carrying it through the airport to save the 13.95 a day. And so because the cheapskates are an adversely selected population relative to the full population, so when I cut my price in this add-on model, I get a bunch of cheapskates. If the marginal profit-- so in the add-on model, there's this big gap between the profits on the marginal consumers and the profits on the average consumers, whereas in the single good product, the marginal and average consumers are exactly the same.

So what we've done is we've just created this gap between marginal consumers and average consumers, and that ends up raising equilibrium prices. Sometimes you've created an adverse selection problem, and that adverse selection problem makes it hard for firms to price compete against other firms, because if you cut prices, you just get cheapskates on whom you lose money renting the cars so cheap. Questions? No? OK.

Leaves me with 10 minutes, and I said I would do this paper very quickly because this really follows very much the methodology Tobias is talking about on Monday of next week. So Miller and Weinberg is a very similar paper to Bresnahan, and it's a much more modern, well-done, frontier quality piece of applied work. The applied question is whether the rise in US beer prices following the 2008 joint venture of Miller and Coors can be explained as a consequence of just static price competition or if we need some degree of cooperation or collusion among the firms to understand what happens.

And so what it does is it estimates a model that assumes that the firms were paying static-- matching with static differentiated product competition before the merger, and then after the merger, MillerCoors and Anheuser-Busch are partially cooperative. And then the partial cooperative model, what it does is-- Bresnahan had three or four models. They have a continuum of models. The continuum of models is assume it's as if each firm is maximizing their own profit plus kappa times the profit of the other firms.

So kappa equals one would be full collusion, kappa equals zero would be static competition, and kappa in between is some degree of partial collusion between the firms. And then what they note is that in this model, with this continuum of models, you can test whether kappa equals zero to say, does the static competition model explain what's going on or do I need to have some degree of collusion starting after the MillerCoors merger to understand why prices are what they are?

So anyway, I won't go into the details here. But background, I think even graduate students drink craft beer these days, but despite the incredible rise in craft beer, most Americans still drink really cheap beer. So most Americans drink Bud Light, Coors Light, Miller Lite. At least most sales because people buy Bud Light in 30 packs. But there's been tremendous consolidation of beer ownership in the past 20 years globally.

So basically, in every country around the world, the same three companies sell all the beer. And so the same companies that sell Anheuser-Busch in the United States are also the ones that sell Beck's in Germany and sell Stella Artois in England and whatever. But these companies have put together portfolios of beers.

Prior to the 2008 MillerCoors joint venture, Anheuser-Busch, InBev had a market share of about 35%, Miller had a market share of 18%, Coors had a market share of 11%. And again, those are market shares in the whole beer market. If you take away the premium beers, there are many submarkets in which those firms have much higher shares. And despite the fact that these were the overall market shares, Miller and Coors were, in 2008, allowed to form a joint venture and sell their beer together.

So MillerCoors argued that, yes, it looks like the second and third biggest firms in the market are merging together to create a three firm competition to two firm competition, but this is going to be good for consumers because there's going to be enough cost savings that we're going to have an incentive to lower our prices. Particularly, Coors was mostly brewed in Colorado and then had to be shipped all across the United States, which is expensive.

They're like, if only we could produce all the Coors beer in all those Miller plants that are distributed around the country and then share trucks shipping the Miller and Coors to all the beer distributors, then we'd save on the trucking costs, we'd save on the shipping costs. Costs will go down so much that consumers will be better off. And the justice department approved the merger. Anyway, you know, very beautiful graph just giving the basic thing of, did what the firms tell the justice department come true?

Here are the graphs of the prices for Miller Lite, Bud Light, Coors Light leading up to the time of the merger. And you can see beer prices were trending down. After the merger, there was this temporary sale on Miller Lite. But then by a couple months after the merger, prices went up there. So there's just a very clear departure from trend, and the prices of Bud Light, Miller Lite, Coors Light jump after the MillerCoors merger. Their competitor's prices didn't do so much.

So the two largest brands not sold by Anheuser-Busch, InBev, or MillerCoors are Heineken and Corona. Heineken and Corona prices, you don't notice-- maybe you look and see there's a little jump up here, but it seems like they're mostly on trend. So if we take that as a differences in differences thing to say there doesn't seem to be a difference in the cost of producing beer, if there's no difference in the cost of producing beer, what you get is that the 12 packs of the flagship brands went up by about 10%.

Overall, prices went up by about 5%. Basically it's as if they-- if a 12 pack of Bud Light Costs $10, it's like they went from 10 to 11. The 20 packs, it seems like-- the 30 packs, they seem like they went from 18 to 19. But it's basically like prices go up by 5% on average, but 10% on the 12 pack of the standard flagship beers.

So why does this happen? Again, there are three things going on here. There are-- costs are presumably going down because they're now brewing Coors all over the country. But MillerCoors has this unilateral incentive to raise its prices because it knows that if someone goes in to buy Coors Light and the price is now higher than they expect, they may switch to Miller Lite, and that's not a bad thing anymore. And then there's this possible collusion, third effect.

So anyway, what do they do? It's random coefficients, logit demand. Tobias will tell you more about what that means. But basically, your utility from buying a brand j of beer if you live in region r at time t is characteristics of the beers times beta minus a disutility from paying for the beers plus a bunch of fixed effect. The brand fixed effect, the time fixed effect, an unobserved quality of the beer, and this idiosyncratic epsilon like in the Hotelling model. Some people like Coors relative to Budweiser just taste wise.

What's different from the standard logit model is it's a random coefficients where the different people have different preferences over the characteristics, and some people are more price sensitive and some people are less price sensitive. So the random coefficients mean that not everyone has the same coefficients. Some people like hoppiness and some people don't like hoppiness. Some people dislike paying money. Other people don't care so much about how much they're paying.

OK. They estimate this model. They have some instruments for price I will skip over, but they find a bunch of charts like this. So this chart says these are price elasticities that says if you raise the price of Budweiser by 1%, sales of Budweiser go down by 4.4%, or sales of Bud Light. But then people will buy other things. And so in particular, Budweiser sales will go up by 0.3%, Coors sales will go up by over 0.316%, Coors Light sales will go up by 0.351%, Corona sales will go up by 2.279%.

We have this sort of substitution matrix they've estimated. And then they say it's a two-step estimation. First, we estimate demand to find all the elasticities. Then we say what would happen under the different models of behavior. And let me see. Two things I want to say. One thing they find is that they estimate which kappa model best fits the data, and they find that the post-period kappas of like 0.25 to 0.37 best fit the data. So it's as if the firms Budweiser and Miller are putting weight on each other's prices and starting to price higher than seems optimal after the merger.

And here's the counterfactual. So this is what beer prices of-- so 12 packs were, like, $9 or whatever before the merger. So we have this long-term trend of declining prices of 12 packs of cheap beer. They think absent the merger, it would have continued along this path with some rises and falls. I don't know if that's due to costs. But then we have three things happening. The first is this unilateral incentive of MillerCoors to raise its prices because it knows that customers it loses when it raises the price of Coors go to Miller, and that's no longer a bad thing. That would cause prices to go up for this much.

But then there were savings from Coors being cheaper to produce, and they say that the Coors savings would have been enough to offset-- that Coors was not lying before the merger when they said that our costs are going to go down by enough that it's going to offset our incentive to raise prices, and they would say the cost savings are this big. So that would have caused this to happen.

However, the third effect is that it looks like they're now coordinating their pricing, and that increase in collusion is enough to go from these triangles to this line up here. So they're arguing is that there are three roughly equally sized effect. The MillerCoors internalizing, it's doing multi-product monopoly pricing, the cost savings on producing Coors, and the Anheuser-Busch Miller cooperative behavior. All roughly equal sizes but combining together to cause this increase in the 12 packs.

Anyway, I think, to appreciate-- the paper has some strengths and weaknesses or whatever. To appreciate those it's easier after Tobias's Monday lecture, so I figured I would just skip those and maybe you can go back and reread the paper after Tobias's talk next Monday and think about what do we think of the various aspects of the estimation strategy. But it is a Bresnahan-like idea of estimate multiple models, see which ones fit best, and that's how I determine how firms are behaving.

OK. Any questions? OK, then we are done. I guess I get a week off. Tobias is going to be here. Gandhi and Nevo is published in whatever handbook of IO that's expensive and hard to find. But if you look at the same exact document as NBER working paper 29257, so if you look at that version, that's easy to find.