FTC - 15th Annual Microeconomics Conference Day 1 - November 3, 2022

Mike Vita:

Good morning. Good morning, everybody. I think we should get started. We delayed a few minutes to let people get in, but I think we better try to stay on schedule. Some of you know me. My name is Mike Vita. I'm the deputy director for research here at the FTCs Bureau of Economics, and I want to thank you all for coming. This as our 15th annual micro conference. Our goal in this conference is always to combine cutting-edge academic research with discussions of real world policy problems. As always, we're grateful to the Tobin Center at Yale University, for their continued co-sponsorship of this conference. For those of you who are not completely familiar with us at the FTC, just a little few words about what we do as an agency. We're an independent agency that, along with the Department of Justice, enforces the antitrust laws.

In addition, we have a second mission. Our other major mission here at the FTC is Enforcement of Federal Consumer Protection Laws. So we have two important missions. Both of these missions are supported by the FTCs Bureau of Economics, which is a group of about, we have about 80 PhDs right now, which makes it one of the largest groups of applied microeconomists in the federal government. And at this point, I want to add that we are on the market again this year. We hired 10 people last year. We had a great hiring year. So if you're on the market or your colleagues or your students, please let them know that we have a vacancy announcement. You can go to the usajobs.gov, and that's where you would apply. So we have these two enforcement missions, and we believe that the two enforcement missions reinforce and compliment each other.

Competition is most effective when consumers are well informed and can make informed decisions. Consumer protection works best when consumers have real alternatives. And so today's conference, like its predecessor conferences, helps to ensure that the actions that we take as an agency are informed and guided by the best possible economic analysis. So before the first panel of today starts, I just want to make a few acknowledgements and a few official announcements. First, let me take a moment to thank Tom Kotch, Will Violet and Stephanie Aaron of the Bureau of Economics for putting this conference together. It's a lot of work. And then they were supported by, of course, a large group of people within BE who I can't name them all, who helped screen all the great submissions that we got for people who wanted to present today. And of course, as we always do, we have a scientific committee of distinguished academics. This year, the scientific committee was Dirk Bergemann of Yale, Julie Holland Mortimer of UVA, and Catherine Tucker of MIT.

So thank you, the three of you, for serving in that role. It's really important. I also want to thank our great administrative team here in the Bureau of Economics who always do incredible work behind the scenes to make sure the conference comes off seamlessly. And that's Maria Villaflor, Kevin Richardson, Constance Harrison, Priscilla Thompson, and Tammy John. And to our research analysts who step in to help organize the logistics of the actual day of the conference. Jen Snyder, Ken Rios, Jesse Justice, Chris Harris, Chris Carmen, Scott Simson, Maryland McNaughton. So just a few administrative things. If you haven't already silenced your phone, please do so. Please be aware, this is the Constitution Center building. If you leave the building for any reason, you'll have to come back through security like you did when you got in this morning. For those of you who are visitors, you've got a lanyard with a badge.

When you're done with that, at the end of the day, there'll be a box outside where you can drop that off so we can recycle them. If an emergency arises that requires you to leave the conference center but remain in the building, just follow; there'll be instructions given over the PA system. If there's an emergency occurring that requires evacuation, then a big alarm will go off. Just follow everybody else out of the building. Importantly, the restrooms are when you go out the door and straight across, that's where those are. And lastly, please be advised, this event could be photographed, webcast, or recorded. So by participating in the event, you are agreeing that your image and anything you say or submit may be posted indefinitely at ftc.gov, or any of the commissions publicly available social media sites. So choose your words carefully. Okay, I think that's all I have. So I'm going to turn it over to Tom Kotch, who will introduce the first session. Thanks very much.

Doug Smith:

Hi, I'm Doug Smith, not Tom Kotch, but I'm also an economist here at the FTC. And our first paper session of the day is organized by Dirk Bergmann. And our first paper will be by Zi Yang Kang from Stanford, who's presenting on robust bounds for welfare analysis. So please give him a warm welcome.

Zi Yang Kang:

Great. Is the slides good? Does the clicker work? Let's see. Great. Thank you so much. So good morning everyone. Thank you so much to the organizers for inviting our paper. This paper is about robust bounds for welfare analysis. My name is Zi Yang Kang. This paper is co-authored with Shoshana Vasserman. So in economics, many papers that quantify welfare changes have the following structure. So first, these papers consider a policy change such as a tax or subsidy that was implemented in some market. Second, these papers then use prices and quantities before and after the policy change in order to estimate something about demand. And finally, these papers use these demand estimates in order to impute the changes in welfare. So the starting point for our paper is that measuring welfare often requires taking a stance on what the demand curve looks like at unobserved points. And so this requires interpolation between the prices and quantities that we observe along the demand curve.

Ane common approach that papers take in order to do this sort of interpolation is to often assume, for the sake of convenience, functional form assumptions in order to interpolate between these two price-quantity pairs. And so common examples of these functional forms include the constant elasticity of substitution or a linear demand curves. And so one really nice example of this is this paper by Amiti, Redding, and Weinstein, published in 2019 in the Journal of Economic Perspectives. And in this paper, they evaluate the change in dead weight loss due to the Trump tariffs in 2018. And so by treating each product month as a separate market, they're able to measure these price-quantity pairs before and after the introduction of these tariffs. And they linearly interpolate between these two price-quantity pairs in order to estimate the dead weight loss as the area below that demand curve. And that's equal to the area of the shaded region.

And so the point here is that this linear interpolation between these two price-quantity pairs is not entirely without loss of generality. And so here I've plotted four different kinds of interpolation that you might possibly imagine. The first is a linear interpolation, but there's also an exponential interpolation, a constant marginal revenue interpolation, constant elasticity of substitution interpolation. And so if you look at the left hand side of the graph, this is at the start of the trade war. This is when the tariffs implemented were relatively small. And so all the different interpolations yield roughly the same estimate of the dead weight loss. However, on the right of the graph, this is when the trade war was in full-blown effect, tariffs were about 30% to 50%. And so if you look at the difference, say, between the linear interpolation and the constant marginal revenue interpolation, this is a difference of $1 billion per month.

And so the point here is that the choice of functional form assumption that we make can have potentially large impact on the sort of dead weight loss estimate that we obtain. A way that the literature takes to avoid making these functional form assumptions is to have a more conservative assumption, conservative type of interpolation between these two price-quantity pairs. An example of this is a really elegant paper by Ge, Segar and Jaeger, published in 2021. This was in the Journal of Public Economics. And in this paper they estimated carefully designed regression discontinuity design to evaluate the willingness to pay or the change in consumer surplus in 1911, due to this pension scheme introduced by the UK government created for poor individuals above the age of 70. And so here they're able to plot the price-quantity pairs before and after the introduction of the pension scheme. And instead of linearly interpolating between these two points, now they have a conservative interpolation between these two points. This gives us the smallest possible estimate of the willingness to pay or change in consumer surplus due to the introduction of the pension scheme.

And so in this paper, our question, we want to compliment these two approaches. And so our question is, is there a more principled way to engage with these assumptions and say something about the change in welfare? And to answer this question, instead of interpolating between these two price-quantity pairs, we establish welfare bounds instead. Our welfare bounds are robust in the sense that they give us the best case and worst case estimates that are consistent with a broad set of pre-specified economic assumptions, which I'll get to in a couple of slides.

And these bounds that we obtain are also simple in the sense that we can compute them in close form. So you want to pitch this as a tool for empirical microeconomists. And so, to be very clear, our bounds apply directly to settings with exogenous policy shocks similar to the tariffs example or the pensions example that I gave earlier. And these are settings where we have measurements of prices and quantities before and after the policy shock, although I'll get to it later in the talk, but we can relax these assumptions. And finally, these are precisely settings where we are interested in effects and consumer surplus or other welfare measures such as dead weight loss.

In the paper we show how our bounds can be applied to a wide range of settings. And so in addition to the dead weight loss on import tariffs or willingness to pay for old age pensions act, which I've talked about, we also have extensions or applications to the welfare impact of energy subsidies as well as the marginal excess burden of income taxation. Of course, this being a 30 minute talk, I don't expect to get to all these applications, but this is just to say that in the paper we envision the set of tools being applicable to a wide range of problems.

So the basic model, the most basic version of the model that we have in the paper builds directly on the two examples that I've already given you. And so just to be very clear, there are these two price-quantity pairs that we observe that lie on an unobserved demand curve. And so the question is what is the change in consumer surplus between these two price-quantity pairs? And of course, if we could observe the demand curve, as any student of economics would be able to tell you, the change in consumer surplus is simply equal to the integral or the area under the demand curve. But here, of course, the main challenge is we don't observe the demand cuff. And so this requires us to bound the change in consumer surplus, which is equal to the sum of the areas A and B. Now for area A, we can actually infer directly compute area A just based on the price-quantity pairs alone.

But area B requires some knowledge about the demand curve. And so the question here is really to bound what area B is. The simplest possible approach is to make no additional assumption. So in the absence of any assumption, we just use the fact that the demand curve is downward sloping that already allows us to say something about the upper bound. So here is an upper bound which is flat and then decreases discontinuously at the quantity Q zero. This gives the maximum possible change in consumer surplus that's consistent with a downward sloping demand curve. And of course, you can see how this logic extends to a lower bound that's given by the green curve that you see in the picture. And of course, here you might have a complaint and say this requires sort of thinking that the demand curves can actually be flat in regions and drop discontinuously.

In other words, this requires you to think that elasticities are equal to zero negative infinity. And in practice we have good reason to believe that in most cases demands are rarely perfectly in elastic or perfectly elastic. And so in this paper, the most basic version of the model in this paper, we're going to take this seriously. And so we're going to do this by imposing the assumption that elasticities need to lie inside this elasticity band, right? So this means that elasticities in between these two price-quantity pairs, these are going to lie between epsilon-lower bar and epsilon-upper bar. And so the question here is what is the change in consumer surplus, given this restriction that elasticity lies within this elasticity band and the demand curve passes through these two price-quantity points.

In order to tell you the answer, I need to define one piece of intuitive terminology. That's one piece and two piece interpolations. And so on the left hand side of the slide, you see what a one-piece interpolation is. Here I've plotted a one piece linear interpolation. It simply means I linearly interpolate between these two price-quantity pairs using a linear functional form assumption. And you can extend this to the right hand side of the slide. These are two piece linear interpolations where you linearly interpolate from the first price-quantity pair to an auxiliary point and then from the auxiliary point to the second price-quantity pair. And so this piece of terminology allows me to state what the welfare bounds for the most basic version of the model in the paper is.

That is the upper and lower bounds for the change in consumer surplus must be obtained by two-piece constant elasticity of substitution interpolations. So as you can see in the diagram at the bottom of the slide, this means that you interpolate using a CS demand curve from the first price-quantity pair to an auxiliary point and then to the second price-quantity pair. And of course, by taking the limits where the elasticities go to zero negative infinity, respectively, this gives the initial box bounds that would've obtained if you had no additional information about elasticity.

And so at this point, let me give you a geometric intuition for why this must be true. And so on the left hand side you see the initial configuration that we have, these are the two price-quantity pairs that we observe. We do not observe the demand curve that connects them. And on the right hand side of the slide, I've plotted exactly the same two points except this is log-price, log-quantity access instead of the price-quantity access, right? So the first piece of information I want to use is the fact that elasticities must be lower than epsilon-upper bar. So this means that the demand curve must be less in elastic than epsilon upper bar. And because this is a log-price log-quantity access, the straight line that you see that is the boundary of the blue region that is exactly corresponding to the elasticity epsilon upper bar that passes through the first price-quantity pair.

And so what it means is that any feasible demand cuff needs to lie on the right hand side of the blue shaded region, right? And of course, this is only one side of the elasticity band restriction. You have the other side of the elasticity band restriction. So the feasible demand cuff needs to lie below the orange-shaded region. And you can repeat this argument for the other price-quantity pair, meaning that the feasible set of demand curves need to lie in this trapezium unshaded region in the middle. And so the extreme demand curve, which would give you the upper bound and the lower bound, these would lie on the boundary of this parallel, this trapezium-unshaded region. But remember that this is the log price log quantity access. And so we need to transform everything back onto the original price-quantity access, and that gives us these interpolations.

Now notice that straight lines in the log price log quantity access, these correspond to curves with the constantly elasticity. And so this means that the interpolation that we see on the left hand side of the screen, these must be two-piece constant elasticity of substitution interpolations. Now at this point, one question you might have is where do these elasticity bands come from? And so even before you think of anything that's on the slide, in the ideal case scenario, these might come from the empirical work that they've already done in order to make these price-quantity pair measurements. And so an example is a regression discontinuity design where you sort of already have randomization around these price-quantity pairs and you know something about the local elasticities at these two points. You might be able to infer something about the elasticity in between these two points.

Without those estimates, you can also use estimates from the literature. So this is an example where we can take elasticity measurements from the literature, say, for short run gasoline demand and say that this is between minus 0.2 and minus 0.4. These give you candidate values for these elasticity band assumptions, or you could also draw from institutional knowledge about the product or related products. So you know that at the extreme elasticities cannot be possibly lower than minus five, and that gives you a plausible bound. A third alternative that I'll talk about later in this talk is that you could sort of draw a symmetric band around the average elasticity. So this is not to say that you know exactly what epsilon-lower bar and epsilon-upper bar are. This is to sort of parametrize the uncertainty or the amount of variation that you allow in terms of the average elasticity or the elasticity in between these two price-quantity pairs that you observe on the demand curve.

And so the basic version of our model requires strong modeling choices. And here I list four of these modeling choices. So first, we assume that nothing is known about the curvature of the demand curve, but in reality, sort of the models that we estimate we take from theoretical models and theoretical models often already make some assumptions about demand curvature. A second modeling assumption that we make is that both of these price-quantity pairs are observed. In reality, one of these points might be a counterfactual point. So you might only observe one point in reality, but you want to extrapolate from this point to a counterfactual. Of course, you could go in the opposite direction and say you could have more data points than we allow for.

And so here we only have two data points. And fourth, the price-quantity pairs that we see on the demand curve we assume are observed precisely, but in reality you might think there might be some measurement error or sampling error involved in the measurement of these points. And so for this next part of the talk, I want to show you how the basic results that we have for the basic version of the model extend in these four different directions. And because this is a short talk, I don't expect to show all four of these extensions. And so for now, let me focus on the first two.

So the first extension is about demand curvature. And so we show how demand curvature assumptions lead to tighter bounds. So where do these demand curvature come from? So on this slide I have this quote from Paul Krugman in his seminal paper on international trade, where he says that the results that he gets in his paper depend on the fact that the elasticity of demand falls with price. And he claims that these results depend on this assumption that elasticity of demand falls with price. And of course, Krugman is actually not the first to come up with this assumption that elasticity decreases with price. This assumption dates all the way back to Alfred Marshall and his principles of economics in 1890. And so since then, this has become known as Marshall Second Law. And there's been an empirical literature verifying that this assumption holds for a broad range of settings of empirical interest.

And so we scoured the literature for further demand curvature assumptions, and we came up with a few more. So, for example, we have decreasing marginal revenue. We also have log concave demand, concave demand, role concave demand, which generalizes log concave, and concave demands. And so we call all of these assumptions concave-like assumptions on demand. And once you think about concave-like assumptions on demand, you can also think about convex-like assumptions on demand. This includes convex demand, law convex demand, and role convex demand. And so these assumptions that I've talked about concave-like assumptions and convex-like assumptions, these are not unrelated to each other. And so in the slide I show you the relationships between these different types of assumptions. So, for instance, concave demand implies log concave demand. Log concave demand implies decreasing elasticity and decreasing marginal revenue, whereas log-convex demand implies convex demand.

And so using these concave-like and convex-like assumptions, we can extend our bounds to include these assumptions on demand curvature. So this result says that the lower bound for the change in consumer surplus, these are attained at the one piece interpolations for each of these assumptions. The type of interpolation depends on the assumption. So for instance, the lower bound for the change in consumer surplus, assuming decreasing elasticity is given by one piece CS interpolation. What I don't show on this slide is the upper bound. The upper bound is also easily stated, the upper bound is actually the two piece interpolation version of each of these interpolations. So, for instance, the upper bound for the case of decreasing elasticity, that's given by a two piece CS interpolation. And these are for concave-like assumptions. And so for convex-like assumptions, all we need to do is we need to flip the direction. And so the upper bound for each of these convex-like assumptions, these are given by the one piece interpolations and the lower bound for the convex-like assumptions. These are given by the two piece interpolations.

Now these results already shed some new light on a previous graph that I've shown you. So one of the things that you might have noticed the first time I showed you this graph is that all these different shapes actually follow a certain order. And this is not entirely coincidental, as the result suggests. And so, for instance, here, the linear interpolation is always higher than the exponential interpolation. And you can interpret this in light of the previous result where I said that the convex demand and log-convex demands are related. And so log-convex demand, actually implies convex demand. And so the upper bound for log-convex demand this should be lower than the upper bound for the convex demand. And this suggests that linear interpolation should always be higher than exponential interpolations. A second takeaway that you can get out of this graph given are new results is that these interpolations can actually be interpreted as bounds themselves.

So what do I mean by that? The upper bound for a convex demand curve that's given by a linear interpolation. And so if you thought that demand was convex, then all the possible estimates for dead weight loss need to lie below the linear interpolation, which is the circle that you see in each column. And if you took Krugman's assumptions seriously and you thought that demand had decreasing elasticity, this means that the constant elasticity of substitution forms a lower bound for that. And so if you thought that demand was both convex and satisfies decreasing elasticity, this means that the feasible dead weight loss needs to lie between the circle and the horizontal line in each column.

So that summarizes what I wanted to tell you about demand curvature. Let me talk about the second point, which is how to extrapolate from fewer than two possible price-quantity pairs. All right, so this is a setting where we observe only one point along the demand curve and we need to extrapolate this to a counterfactual point. So here I've illustrated a counterfactual price called that P one, and the question that we want to ask is what is the change in consumer surplus when I increase price from P zero to P one? And here the counterfactual price is something we already have, but we don't observe the but for quantity. And so that's why we only observe one point along demand curve and not the second. And so the nice thing about this extension is that all the geometric intuition I showed you from before actually extends to this case. And so in fact, for this case, the upper bound is given by both the upper bound and lower bounds, in fact are given by a one-piece constant elasticity of substitution interpolation. And that's shown in the graph on the left hand side

Now, so in the paper, we apply these results to the CARE program. The CARE program stands for California Alternative Rates for Energy. This is a program that provides a 20% discount on gas for low income households. And what this entails for low income households is that given the subsidy of 20%, gas consumption actually increases, which entails an increase in consumer surplus, but this comes at a cost to the environment because of marginal social costs of gas consumption. On the other hand, there are also households who are not eligible for the CARE program. These are what we refer to as non-CARE households. Gas prices increase in equilibrium for these households because we had to provide a subsidy for CARE households. Prices for everyone goes up. This means that gas consumption of non-CARE households go down, consumer surplus for them goes down. And of course, this entails some benefit to the environment. And finally, there's a third component of this buffer analysis, which is that the CARE program is expensive to administer. This entails an administrative cost of $7 million.

And so conceptually, this is very similar to sort of the graphs that I've been showing you throughout this talk. So on the left hand side there is the demand curve for CARE households. On the right hand side, there is the demand curve for non-CARE households. The one thing I want to draw your attention to is that the counterfactual that we are considering is what is the counterfactual when there is no CARE program. And so in the absence of the CARE program, what happens is that there's uniform pricing across the two different groups of people. And so this uniform price is denoted by P-star on this slide. And so for the CARE households, prices go down because of CARE. So instead of P-star as a uniform price, they face a lower price now given the energy subsidy. And so this entails a sort of net welfare impact on CARE households given by the green shaded region. Whereas with non-CARE households, prices actually go up from the uniform price of P-star. And so this entails a welfare loss equal to the rich shaded region shown on the graph.

And so given this trade-off, the question that we want to ask is whether the CARE program is net welfare improving or not? And so this was actually done in a really beautiful paper by Bob Hahn and Rob Metcalfe. This was published in 2021 in the American Economic Review, where they had this empirical strategy of randomly nudging eligible households to sign up for the CARE program. Using that, they compute the late-on average treatment effect based on gas usage. And they interpret this treatment effect as in elasticity with which they used to answer the question, how much does gas usage given, this 20% discount, how much does gas usage change given this 20% discount per unit gas prices?

Their empirical exercise requires a number of modeling assumptions. One assumption is that this counterfactual uniform price P-star, this is given by this accounting equation over here. This just means that the uniform price P-star can be obtained from these price-quantity pairs from CARE households and non-CARE households. They know the counterfactual price P-star, but again, they do not observe the counterfactual quantities at P-star. And one other assumption that they require is that consumer demand is linear, right? So here they assume that both the demand curve for CARE households as well as the demand curve for non-CARE households are both linear demand systems.

So they estimate a CARE elasticity of minus 0.35. They obtain this elasticity of non-CARE households from another paper in the literature. This is a paper by Alf Hammerer and Ruben, and they obtained this elasticity to be negative 0.14. And so these elasticities imply that for CARE households, the net welfare gain is actually $5.3 million. For non-CARE households, the net welfare loss is negative $3.1 million and net of administrative costs of $7 million. There's an overall welfare loss of $4.8 million. Now in their paper, they have a bunch of robustness exercises around this result. And so using these robustness exercises, they ask the question, what would the net welfare impact be if the elasticity was actually not actually minus 0.35 for CARE households? Crucially, these exercises still require the assumption that demand is linear. And so this is exactly where our results of this paper come in. So instead of assuming that demand is linear, here we show what happens if demand was not linear.

Each point in this graph corresponds to an elasticity band first around the CARE elasticity of minus 0.35 as well as an elasticity band around a non-CARE elasticity of minus 0.14. And so at the lower left hand corner of this graph, this is precisely the case where you have no uncertainty at all about variation in elasticity. This is when for sure that elasticity in between these two points are given by minus 0.35 and minus 0.14 respectively. Contrast this to the upper right hand corner of this graph, this is when you have a lot of uncertainty you are allowing for a lot of variation and elasticity in between these two price-quantity pairs. So there are two takeaways from this graph. The first takeaway is that there's a lot more red than green, as you can see. This does not require the assumption that demand is linear.

And so this suggests that even without that assumption, even dropping that functional form assumption that how did Metcalfe make the result that the net impact of the CARE program is negative seems to be very robust to these different functional form assumptions. In other words, you need to allow for a lot of variation in elasticity in between these two price-quantity pairs in order to flip the result.

Now, a second takeaway from this graph is you might be thinking this result, it uses the fact that administrative costs are $7 million and that's a big amount. And you might ask what happens when instead of $7 million, it looks for like $2 million, right? So if that were the case, then the amount of red and green in this graph would be sort of more evenly distributed, and I think that is exactly where our result has the most bite. This allows us to of parametrize the amount of variation in elasticity that we allow and sort of say how much variation in elasticity must we allow in order for the result that the CARE program has a net welfare loss to flip and say that the CARE program has a net welfare gain. Okay, so in the paper we show how the results extend also to the third and fourth point.

Zi Yang Kang:

... points, so we interpolate with more observations instead of just two, and we show how to incorporate sampling error into our welfare bounds. Beyond these four extensions in the paper, we also have further extensions. So for instance, we talk about how this works with producer surplus, we can handle more heterogeneity, answer distributional questions, can handle alternative welfare measures, and handle multiple objectives at once. And so I'm running out of time so let me conclude. In this paper, we develop a framework to bound welfare based on economic reasoning. We hope to make the case that everyone should use this. We draw, assess conclusions from empirical objects that are commonly estimated and we're very excited about this. I look forward to the discussion and thank you very much for your attention.

Doug Smith:

Thank you very much, Zi Yang. So we'll now have Christopher Adams from the Congressional Budget Office as a discussion for the paper.

Christopher Adams:

Doug, can I have some [inaudible 00:36:39] Thank you. First, thanks to Doug and Dirk for inviting me to give this, and thanks particularly to Tom, Will and Stephanie for carrying the baton and getting us to the

15th annual conference. I love this paper. For me it sort of hits sort of three points that I care about. I'm a policy analyst at CBO, so I like when academics provide us with tools. I know there's a number of analysts here. Of course at CBO we don't care about consumer welfare. I forgot, does the FDC still care? But let's assume that somebody cares about consumer welfare, then they have a really great tool.

Second, it's bounds. I've loved the approach that Chuck Mansky has been promoting since the early 90s about using bounds, and I particularly like the idea that John Pepper and Chuck Mansky have been promoting where as analysts, we present the results in a order, from the most credible assumptions to the least credible assumptions, and allow the policy makers to see how the assumptions are driving the results. Lastly, it's demand estimation. Who doesn't love demand estimation?

So this is the paper. The basic setup is that we have these two points. We have a pre-policy point, P0Q0. We have a post policy point of P1Q1. We want to determine what A and B is. A is fairly straightforward to determine, B we're we're less sure of. So just sort of going through this approach of going from most credible assumptions through to the least credible assumptions, I was trying to think what do we know if we just know these two points? Potentially, we know that they come from a demand curve. Through this here I'm going to present results. They're not really results, they're just my thoughts.

Mansky called this worst case bounds, Pearl calls his natural bounds, probably doesn't say anything. So we have to make an assumption, we have to make a structural assumption in order to move forward here. So the authors used demand slopes down. I would've liked a bit of discussion about the stuff that we had to go through in grad school on GARP and WARP. It'd be nice to get a little payoff for that. So the assumption, basic economic assumption says, well B must lie in this box from this green line to the red line. Now I've drawn these as being pretty similar. You might have cases where A is pretty much the whole ball game and it really doesn't matter. B is so small relative that it doesn't really matter and so you're sort of done. Once you get A, you're sort of done. And so that's a useful thing to sort realize from just from this setup.

If we are willing to add an additional, stronger assumption like Marshall's second law, and Marshall's second law says if we have a price change, we would expect at high prices that there's going to be a big change in quantity, so our curve's pretty flat. At low prices, we have a price change, we expect the quantity change to be pretty small so we have this concave demand curve. The authors go through a bunch of different assumptions, but the immediate impact of Marshall's second law is to increase the lower bound. So we get tighter bounds. I realize that this picture's probably wrong. I think the green line is wrong.

But anyway, what if we added additional information? We have elasticity estimates, we've got our own, we've looked in the literature. We've done an approach like an empirical basion approach, which gives sort of probability distribution over a set of elasticity estimates. Then once we have these elasticity estimates, then we can get tighter bounds. And this really is the technical contribution of the paper that the sharp bounds here are given by the two CES functions. Just as a side, there was a interesting paper given the Chamberlain lecture last week on how to get sharp bounds and the authors might want to look that look at that just as a different approach to proving their theorems.

So the proof of this theorem turns out to be this picture. When I went through grad school, I was told that Ray Danica would tell us that a picture was not a proof. Turns out he was wrong. Picture is the proof. They have a more complicated theory in which I didn't go through that. But if you want to do that, you can do that. So my interest was a slightly different question was, what if we knew at one point, P1Q1, but we are interested in sort saying something about the whole demand curve? So that's something I'm interested, I know [inaudible 00:43:14] is interested in from her other work. So I wouldn't sort think through that exercise. Again, take the pictures with a grain of salt. But this is just saying, well what if we had some upper bar, upper price and upper quantity? These are what the bounds would be.

More likely, we have some estimate for our elasticities. So this we have of a bunch of elasticity estimates that are in a pretty small area, but we want to of say something about this quite large area. I think we'll need to work out what these endpoints are, what P0 and Q0 are given the assumptions like Marshall's second law. So this straight off, we learned from this paper that this picture is not right, that we can use CES function and Marshall's second law to give us tighter bounds. And then, we can do better than this. Again, I think this picture is wrong, but if we have elasticity estimates from the literature, from our empirical basia model from something else, then we can say more.

So this is somewhere that I'd like to go. We're going to have to work out where P is and Q is, but I think that would be in interesting, at least for me. The paper is very polished. So there's probably the advice is probably don't do any of this, just get the thing published. So I'm going to finish with that. I really enjoyed the paper and I'll hand it over to questions and answers.

Tom Kotch:

So if there are questions from the audience, there are two microphones because we are doing this into the rest of the world. In order for people on the internet to hear us, you do need to speak into a microphone. So if you have questions we will find you. Just raise your hand.

Steve Barry:

Hi, this is Steve Barry from Yale, and I'm cheating because I saw your co-author give this paper on Tuesday and then I thought about it afterwards. So let me just say two quick things. One is, I'd say a little caution about Marshall's second law. It's natural I think to think of it in a context of a representative consumer as the price gets higher, the alternative good looks better and better as a substitute. But in for example, discreet choice modeling, what often happens is the price elastic consumers just drop out as the price rises. You're left with the inelastic consumers, which is why you have this opposite intuition that high price goods are inelastically demanded, right? Because you've only got rich people left to consume them. So the opposite story often comes out of a class of discreet choice models.

And that leads me a little bit to pick up on something that Chris said, which is I think it'd be fun to take this paper back and think about the bounding empirical literature, right? So one idea is you have these theoretical restrictions. Another idea of is that you have two different instruments. They're both bad, but one gives you an upward bias and one gives you a downward bias, something like that, right? And think about cases where you can actually get the bounds out of some kind of set of weaker assumptions say than that you have a perfect instrument.

Zi Yang Kang:

Great, thanks very much. I'm also cheating by the way, Steve, given that my co-author has conferred with me since talking to you. But the short answer is yeah, we are thinking about these issues and we're working on both of these issues in follow up work.

Tom Kotch:

So thanks everyone in the audience for listening to the great presentation and discussion. And Doug, if you move us on to the next paper.

Zi Yang Kang:

Thanks.

Doug Smith:

Thanks again to Zi Yang and Chris. So the second paper in the session, will be presented by Kevin R. Williams of Yale University and NBR. The paper is entitled, Dynamic Price Competition: Theory and Evidence from Airline Markets.

Kevin R. Williams:

Okay, good morning everyone. Thank you for being here, and for those of you online, thank you for being there. And thanks to the organizers for including this paper. This is Joint Work with Ali Hortaksu at Chicago, and Anika Ori, my wonderful colleague at Yale. Acknowledgements, we'd like to thank the Tobin Center, [inaudible 00:49:30] for providing resources that supported this project. And to Jose, who is an excellent pre-doc and he will be applying to graduate programs this fall, please look out for him. What I'm about to show you uses confidential data. None of us have any financial relationships with the entities described in this research. Okay. Dynamic pricing is commonly used in markets with fixed capacity in a sales deadline. We're thinking about planes, trains, cruises, entertainment tickets, and some aspects of retailing. And importantly, capacity drives the pricing dynamics.

And that's because the opportunity costs of a seat changes due to scarcity. And the value of this seat or a unit of capacity depends on my ability to sell it in the future. This is where we get the intuition that if I have a bunch of excess inventory, I'm going to fire sale or offer low prices. In addition, demand might change over time. Firms want to supply units to those with the highest willingness to pay. If those consumers arrive close to the deadline, this would give firms the incentive to hold back inventory to supply those consumers. This paper is about competition. What additional forces does competition bring? First, there's competition. So the demand today depends on all prices, not just a single firm's price, firms face their residual demands. And in addition, the opportunity costs of selling not only depends on my units remaining, but also your units remaining precisely because they determine future prices.

For example, if I have excess inventory, I might charge a high price instead of a low price because I want you to sell out. If you have one seat left, maybe I want to get you to sell out, then I'm the only firm in the market for the rest of time. Maybe I want to fire sale even though I have scarcity, few units remaining, because it will soften future price competition. So if these things are true, which they are, we have examples in the paper, that means equilibrium prices and profits lack nice properties.

So one question is, can we get a framework or have a framework in order to understand dynamic price competition generally? And second, if we have this framework, what are the welfare implications of dynamic pricing, or DP, in a market we're all familiar with? That's this paper. So first, we're going to introduce a framework to study dynamic price competition with scarcity and a deadline. We're going to provide a differential equation characterization of equilibrium dynamics. For those of you who know this kind of seminal paper by Lego and Van Risen, we're kind of extending this to oligopoly. We provide insights on equilibrium existence and uniqueness. And we study competitive dynamics.

And with this framework, we are going to tackle the welfare effects of DP in oligopoly. And our context is the airline industry. The punchline of our empirical work is we find the opposite welfare effects of previous work. In our setting, we find that dynamic pricing relative to uniform pricing expands output but lowers welfare. This contrasts with existing work in retailing, ride share, single carrier airline markets where DP is found to improve welfare. And then we're going to examine heuristics that try to approximate kind of airline pricing practices. And the main result here is that relative to DP, these heuristics improve welfare. The DP benchmark model results in the lowest welfare of all the models that we consider in this paper. So let me introduce the model. There are J products offered by F firms. Each firm owns a subset of products. Products are partitioned. Firms are endowed with exogenous initial capacity K.

I think endogenizing capacity is very important. That's not in this paper. Firms face a common perishability date T, and we consider time interval lengths delta, we're going to take the limit. Delta going to zero in our analysis. Within a period firms choose prices. A consumer arrives according to a time varying probability. We're approximating plus on here in discreet time. And then this consumer, if she arrives, she decides to buy one of the available products or choose the outside option and leave forever. Firms observe the entire history of prices and inventories. A bit on the demand model, we assume that consumers are short lived.

So every time period corresponds to a well defined demand function. This contrasts with interesting, important work. In theory, that considers forward looking buyers. And with forward looking buyers, a firm is essentially competing with its future self. We're not doing that in this paper, partly because of empirical evidence. But in our paper, firms are competing against rivals, not their future self. So a consumer arrives and she decides to purchase or not. There's a purchase probability, SJETP. This is a function of P, the demand parameters and the set of available products, A. We impose regularity conditions on the demand functions themselves. Please see the paper. This is like assuming long concavity in the single firm, single product setting. What we can't have is multi-peaked demand. Like discreet random coefficients pretty problematic in our setting because there are multiple equilibria.

Jose has found four equilibria and nine zeros in a stage game. So we're not doing that. But our model extends outside of IIA, including nested logit. So we're going to use nested logit in our empirical application. Okay. Solution concept, we're going to pursue mark off perfect equilibria. The payoff relevant state is capacity at time T and a mark off strategy for product J is just PJTK. I want to start with the single firm setting. It's the easiest and it's nice, and I'm going to show you the nice properties and then say, well none of these properties hold in oligopoly anymore. And I'm also going to start with a single firm case because it's useful for establishing notation. So now firm M owns all the products. I started the talk by saying opportunity cost drive the pricing dynamics. What is the opportunity cost? The opportunity cost omega J. So four product J, is just the difference in continuation profit.

If I go tomorrow and I have all my seats left, versus if I go tomorrow and I sell one unit of J. That's what EJ is. It's a zero vector with a one for the product that is sold. Okay? So this is the opportunity cost that drives everything. I'm not going to refer to this, or we don't refer to these objects as opportunity costs in the paper. We prefer to call them scarcity effects because there are going to be many scarcity effects. Okay. So what is the control problem for firm M? Big pie, the continuation profit is equal to this mess. No, it's not. A mess that's beautiful. The first line is what? The probability of an arrival, the probability given prices that J is sold. If I sell J, I'm rewarded with something. That's called revenue, PJ. But if I sell PJ, I lose something. I lose a unit of J, and that is what the continuation value after a sale is.

So that's well defined. And the second row, what's the second row is okay, no one arrived, no one bought anything. And so time moves forward, but I keep all my K, and that's the continuation profit. We already know that this big pie delta, as a length delta, is well behaved. So if we take the limit of delta going to zero pi MTK solves the following differential equation. And so the time derivative is equal to this. Also a mess. No, very beautiful. This is entirely familiar, it's just expected demand time to markup. But the cost here is the opportunity cost or the scarcity effect of K. Okay? So that's the single firm problem. This problem is really nicely behaved. So the value function is decreasing in time and increasing in capacity. I would like to have more time to sell. I'd like to have more seats to sell. The opportunity costs are decreasing in time and capacity.

If time moves forward, and I don't sell anything, opportunity costs are going down. And I'm not sure if this has been established formally in previous work. If we view the omegas as a stochastic process, they are a [inaudible 00:59:08]. That is average prices go up in this problem because of demand uncertainty. So if you simulate one of these DPs, you might notice that it's always U-shaped. The increasing part is this theorem. And why do they go down? Because one is unique. Once I get to one seat remaining, prices are going down. And so that is another nice property. Unfortunately, none of these hold in oligopoly. So back to oligopoly, I'm going to show you just two firms, two products. So one firm owns one product, the other firm owns a different product. All of our results hold to an arbitrary number of firms and products.

Now we have two big pies because there are two firms, and we have many scarcity effects. Why is that? Because if I sell my own unit of capacity, I'm affected. That's my own scarcity effect. But I am also affected if you sell a unit of capacity precisely because they govern the future price path. So we have own scarcity effects, competitor scarcity effects, and now we have a matrix of scarcity effects. So my own, your own and the crosses. And for many products and many firms, see the appendix and Anika has all the beautiful equations. Big. Okay? So now we have a big omega.

So this is one of the key results of the paper. Assume sufficient conditions on demand and assume that the time horizon is not too long. Then the value function as we take the time interval length to zero converges to the following differential equation. Okay. So the time derivative is equal to this. Not a mess, very elegant. The first thing we've already seen before, it's exactly someone arrives, someone buys my product, I'm rewarded with something, revenue price. And if I sell, I'm affected by my own scarcity. But the second row here is the new term that is, if someone arrives, someone purchases your product. I'm affected due to your scarcity. So we're plugging in a particular price here, it's the equilibrium stage game price. And I'll tell you more about that in a second. But this game is parameterized by the entire matrix of scarcity effects and the demand parameters, theta, and there are boundary conditions and all that.

Just pause for a second, we're still in the theory, but this characterization is wonderful for us because now we can simulate large games in reasonable time. So Jose can simulate a game with 250 to 300 million states in five minutes. Okay, so this I think can be useful for empirical researchers investigating dynamic pricing with perishability and a deadline because it allows us to go to the computer. Okay, I want to emphasize two parts about this proposition. First, is go back to the value function just emphasizing the scarcity affects again. Assume there's some sequence of prices, P. What is the value function in discreet time for firm F?

Again, it's someone arrives, someone buys my product, I'm rewarded with revenue, the price, and we move forward in time and I've subtracted off one unit from my K. There's a typo, this should be J instead of F, but they coincide here in the first row. So this is just the revenue from own sale. But there's a second row with two terms. Let's go to the last part. The last part I basically already showed you, that's nothing happens today and we move forward in time. Both firms have their Ks. The new term is that someone arrives, someone buys from you and I'm affected by your scarcity.

We don't plug in any prices. We plug in the equilibrium prices from the sage game. And so here is the stage game. The sage game is justifying as the difference in continuation profits of going T + delta - T. And so the equation there is also familiar. It's just expected demand times a markup. But again, there's this new component, which is how your scarcity affects me. So this means we can't apply results like [inaudible 01:03:57] because the sage game isn't of that form. This tells us a few other things, or creates headaches, but Anika has solved them. One is we can get multiplicity even with logit. So we need to think about sufficient conditions, see the paper. Another thing is this can be a game of strategic compliments or substitutes, even though I haven't endogenize the capacity choice just because of all the omegas or scarcity effects that affect prices today.

Okay. So please see the paper for more details on this, but let me summarize more theoretical results and then get to the data. So profits are not monotonic in own capacity. Profits are not monotonic in competitor capacity. Profits are neither concave nor convex in capacity. Scarcity effects can be positive or negative. And we have some insights on the pricing dynamics. So one thing is we don't like being symmetric. If we end up in a state where I have 10 seats and you have 10 seats, that's the fiercest competition we have in our model. We want to get out of those states. And so that's where you tend to see the fire sales. So we try to get to as asymmetric states over time. The second thing we can establish is that the largest price effects or jumps occur when the firm with the least number of units remaining sells.

So I said previously that I might offer a fire sale even when I have scarce capacity because it's often future price competition. That softens future price competition the most. And we also have a new markup rule in the paper. We don't use the markup rule, but that might be useful for empirical researchers. If I have demand estimates, I have my Ps, I can uncover some of these scarcity effects using the markup rule. Okay, now we're leaving theory and we're going to apply the theory. We're going to analyze the welfare effects of dynamic pricing in my favorite industry. So airlines, we have some pretty incredible data. We have the Ps and Qs of multiple firms selling flights. Which Ps? Which Qs? We have all the Qs, well and all the Ps. So I see booking counts or we see booking counts for all types of traffic, local flow, people going A to B, people to ABC, people going ACB. We're not going to do the network. That's too complicated.

But we see all that. We don't have the exact itineraries. We see if you book from the airline or from online travel agent, something like that. Which Ps? All the Ps, but we're going to focus on the lowest available economy class ticket. The data identifies firms, flight numbers, departure dates, et cetera. And you're like, how is this possible? Think of this as the Nielsen of airline data, although it's a lot of work to get and combine, but that's the idea. In addition, we have all the click stream data. We observe the exact pricing technology and we observe all sorts of outputs from one of the airline's revenue systems. If you go to my website, there's paper called, Organizational Instruction and Pricing. It talks about this, you can dive in there. We're going to use the search data to think about arrivals. That was the Pusan process that I showed you earlier. There are problems, we only observe a subset of searches. So we're going to scale up and I'm not going to talk about that in this presentation.

Some facts about the data. Here's the markets that we study. So we're going to hone in on duopoly markets where a large fraction of traffic is nonstop local. I want the ideal market is everyone is just going A to B. Welcome to reality never happens. So I want to minimize ABC and I want to minimize ACB. That's kind of our basic selection criteria. But the distribution fairs for the markets that we analyze is similar to all duopoly markets. Basically, we're studying large cities to medium size cities. You probably know some of these markets. We're going to have about 60 directional pairs. In the histogram, I have the distribution of monthly departures that count. On the one extreme head-to-head competition, you and I have one flight each. And on the other extreme, there are 10 flights a day. So that kind of gives you an idea of the markets that we study. And the right graph shows us the selection criteria. I don't like the left tail because it's just more complicated. And so our distribution is shifted to the right, the percentage of nonstop.

The percentage of OD traffic, nonstop is shifted to the right. That was on purpose. Pricing dynamics over time. The left graph shows P. There's a step function because firms often use advanced purchase discounts. If you cross the fence, the probability that you face a price increase is very high. Please see the other paper that talks about this. And on the right hand side, I just have the average booking rate across the different flights for an airline averaged over routes, so this just like the booking rate. And this tells us demand is upward sloping because Q is increasing in P. Just kidding, don't panic. That was a joke.

What is going on here? Demand is becoming more elastic and the number of people interested in travel is going up. So this is kind of useful identification, I guess. And that's what I want to say there. So that's dynamics over time. Here's some facts about across firms. So no, at least for our routes, no competitor consistently sells a larger fraction of the plane. That's the dot plot thing. Dot plots are just load factors for one firm and the other

Kevin R. Williams:

... firm, its competitor and you can see it's all over the place. Sometimes planes sell out, sometimes planes don't sell out. The world is uncertain, okay? But importantly, the route level averages. Sometimes one firm sells more, other times for other routes, the other firms sells more. So there's a lot of heterogeneity. The right plot is the average price difference over time. I guess, let's start with the orange line, because that stands out. It tells us that one firm charges a higher amount, a little bit, well in advance, and then tends to be the firm that is a little bit cheaper closer to the departure date. I don't think that's particularly interesting, but one thing that is interesting is there are two lines, but they coincide, which tells us that 50% of the time you and I charge the same price. Please see the other paper for why that happens.

That's also useful variation that is going to appear in our demand system. So here's the demand system. It's nested logit. There's a product, there's a departure date, there's a time from departure, and there's a route. People care about when they fly. People don't like paying too much for a plane ticket. We allow the price sensitivity to change over time. That's why it's alpha T. And this is nested logit, where there are two nests, the outside good and all inside goods are in their own nest because I care about flight substitution in that way. The sigma captures that. Okay, It's a static discreet choice model. There's no casie, we love casie, put casie in. Okay fine, we can do that too. So there's a casie that's potentially correlated with P. And because persona is exponential, we can estimate this with a control function. It doesn't really affect our parameter estimates. If you allow for an additional unobservable that's correlated with price or demand elasticities are slightly higher than what I'm about to show you.

So we have consumers, they have preferences, but how do consumers get to the market? This is our Phrosan process and it's fairly rich. It depends on the route, it depends on the departure data, it depends on the date people search, it also depends on the number of days to departure. So there's many parameters here. And since we're going to do this as a approximation to the continuous time model, it could be that demand is censored. We have one seat left, but two people demand it in this time interval delta, so we accommodate that and as I said, I'm not going to go into the search scaling up arrivals, but our results are also robust to various hyper parameters on how we do that.

Demand estimates. Demand is downward sloping and demand becomes more inelastic as the departure date approaches, that's the alpha T on the left and on the right we have our estimated or recovered arrival rates. The interest in travel is increasing as the departure date approaches. We estimate that flights are quite substitutable. The sigma is 0.5 and the demand elasticities are around 1.4, which is substantially more elastic than what I found in previous work in single carrier markets. And now we have demand. Let's go back to supply. So we're going to use the differential equation characterization to simulate market outcomes, that's the benchmark model. And then I'm going to shut down dynamic pricing entirely for both firms and have a uniform price.

What to emphasize? Firms prefer dynamic pricing. Actually the revenue drop of uniform, that's pretty substantial, it's 14 to 16%. So firms like dynamic pricing. Consumers do not like dynamic pricing. Consumers prefer uniform pricing, and again, the magnitude is quite large. Our welfare results are opposite of what has been found previously in the empirical literature, to my knowledge. Under dynamic pricing, output goes up, but welfare goes down. What is going on here? Essentially this is a competitive model.

High willingness to pay arrives close to the end. So firms are incentivized to create scarcity early on. There's scarcity early on. It allows them to have higher prices towards the deadline. But then there's a misallocation, like there's too much maybe scarcity being created because the number of sellouts is higher under dynamic pricing than it is under uniform pricing. That's the scarcity effect, of my scarcity and your scarcity. How they affect the time path of prices and yield these welfare results. So I can show you that and emphasize it one more time. The left is product shares, uniform and dynamic pricing. And then the right hand side is the cumulative ratio of dynamic pricing welfare over uniform pricing welfare. So dynamic pricing, much lower prices early on, higher product shares, creating scarcity, drives up prices at the end. And if you go to the right hand side, you can just see that the cumulative ratio is going down a lot and it flips close to the deadline, and that's why or how the result that dynamic pricing lowers welfare, but has higher output than uniform pricing.

So that's dynamic pricing. That's uniform pricing. We consider two heuristics that try to get at actual airline practices. This full DP model, very complicated, lots of moving parts, I don't know of a firm that does something like this. And the other paper we tell you how airlines price and we're trying to get at that with these heuristics. So we consider two. The first is, I'm a firm, I know you exist, but I'm not going to internalize all your omegas, that's too hard. And so what I'm going to do is I see what price you charged yesterday and I think you're going to charge it today. So that's one benchmark model. And then, sorry, it's not benchmark, it's heuristic, counterfactual.

And then the second heuristic we consider is even simpler. It is firms, airlines use a discreet menu. Again, please see the other paper for all the details on this. But there's a menu affairs, I know your menu, you know my menu, and I think you're just going to follow the min of the menu over time. So those are the two heuristics. And let me just emphasize one number or result. Welfare is higher. So these heuristics raise consumer surplus and welfare relative to the benchmark dynamic pricing model. The dynamic pricing benchmark model results in the lowest welfare of all the models that we consider.

And that's the paper. So we hope this paper provides a framework for studying these dynamic pricing with competition. We see this everywhere in hotels, cruises, airlines. So it's a framework to study them. And then we've applied it to an important industry, the airline industry. We show how competitor scarcity affects price dynamics and we find the opposite welfare effects of previous papers, including my own. And there's a lot more to do and we're working on it. That's it. Thanks.

Doug Smith:

Thank you, Kevin. Now, Juan Ortner, from Boston University, will discuss the paper.

Juan Ortner:

Thank you. Oops. Okay. Right. Thank you very much to the organizers for having me and for giving me the opportunity to discuss this pretty impressive paper. So let me just say, I mean the paper is, as I said, impressive. It combines very nice theory with amazing data. So it's a pleasure to read. So let me go briefly over what the paper does. So the goal of the paper is to study dynamic pricing under oligopoly when you have these capacity constraints and deadlines. So in terms of the theory, the main result, as Kevin just said, duopoly is very, very different from monopoly.

Now pricing depends on a complicated way, not only on your capacity but on how you expect the capacity of your competitors to evolve. And prices can be complement, substitutes and a lot of things that we thought or that happen under monopoly don't happen under duopoly. And then on the empirical side, as I said, the authors have this impressive and amazing data and that allows them to estimate the model and perform some counter factors like what would happen if there was uniform pricing rather than dynamic pricing. And well what would happen if firms used some heuristics?

So briefly the model, although Kevin did a great job explaining this. So we have end firms, they sell J products. These are imperfect substitutes, must be sold by some date capital team. At each point in time, firms are going to choose the prices that they control and then consumer arrives randomly and consumers are not forward looking. There's some demand function. This as captures the probability that this consumer purchases good J and firms have some initial capacities. And these capacities it's important, these are the state barrels, they are perfectly observed, each sale reduced them. If there were just one firm, the problem would be quite easy. The firm would choose this monopoly price where they're maximizing profits. There's this endogenous cost which is this omega J, which captures kind of the change in your continuation profits, if you sell a unit of good J. I mean it's still complicated because you have to take into account the sales that you make in each of these routes.

So for an airline that has thousands of routes all over the country or all over the world, even keeping track of this is complicated, but can be solved, and the monopoly model is really tractable. Profits converge when the timing interval goes to zero, there's this well-behaved ODE that characterizes profits so you can back out prices and everything. So the monopoly model is pretty simple to describe and to characterize. The duopoly model, things are more complicated because now as Kevin said, when you make a sale or when you choose your price, you're not only considering the effect that that price is going to have on the chances that you make a sale, but your price might affect the chances that your opponent makes a sale. And if your opponent makes a sale that's going to affect their capacities and that might be good or bad for you.

So this omegas, now you have to keep track of well what's the effect that your opponent making a sale is going to have on your profits tomorrow? And if your opponent makes a sale that might lead to lower prices or higher prices, might benefit you or might not benefit you. So prices can be compliments or substitutes, you can have multiple equilibrium but under some conditions, there's going to be a unique algorithm, and again you can get ODE that's going to characterize profits.

And then in the empirical analysis, basically they estimate consumers arrival process using the booking data they have and that allows them to recover prices through this ODE that they characterize the theory part. They have these two counterfactuals. One is what would happen under uniform pricing? What would happen if they did this pricing heuristics. Uniform pricing leads to higher consumer welfare and higher consumer surplus, then dynamic pricing leads to lower revenue. And it's driven by this effect that they at this in the uniform pricing case, those consumers who arrive later now, that they have high valuation, they're getting the good at a low price.

The consumer, this pricing heuristics also lead to higher welfare, higher consumer surplus. And one of these heuristics also leads to higher pricing, which I was talking to Kevin this morning. This makes you wonder whether if these heuristics kind of are close to what airlines actually do in practice, this makes you wonder whether this might be some collusive algorithm that they're using to obtain higher price, higher profits relative to what they would get in akili from if they... Anyway, the pricing heuristics that the paper considers have the following features. So basically this pricing heuristics take, or basically what they do, I don't know if I can go back, they kill these kind of strategic effect that you want to change your price to induce your competitor to sell more. So both of those heuristics do that. And when I was reading this that kind of made me think or reminded me of the solution concept of oblivious equilibrium in which state barrels are too complicated and you cannot keep track of everything.

Here is kind of the same. You're not keeping track of the effect that your prices are going to have on your opponent's capacities. So that made me think, well how could you adapt this solution concept or think of something related to of oblivious equilibrium in this setting? And one thing that you could think of is, well maybe firms take us given how your opponents capacities are going to evolve, they think of those as evolving exogenously and they're just maximizing profits taking into account their own capacities, which would be that pricing equation that I have there. So I mean that seems something tractable to solve. And this also made me wonder, the authors talk about the equilibrium as the benchmark, but that doesn't seem to be the benchmark, that's not really what the firms are doing. So I guess, would be nice to know which of these pricing rules explains the data better. Whether it's this one or the heuristics or some other one.

Kevin talked about this, I don't know that it's in the paper, but when they estimate demand, they assume that there's a lot of parameters that are constant across routes. I don't think it's in the paper. It would be nice to know how heterogeneous these routes are. Are they all between large cities and small cities? Are they all on weekdays versus weekends? Because you might imagine that many of these parameters would be different for different types of routes or for different days of the week. Anyway, I don't think it's in the paper, it would be good to know a little bit more about that and maybe re-estimate them all separately for different types of routes to see whether that makes a difference.

One of the results that the authors highly emphasize is this thing that in contrast to previous studies, uniform pricing leads to higher consumer surplus and higher welfare than dynamic pricing. And that's true in their model. But it's also true that this uniform pricing leads to much lower revenues for the firms. And now that makes you wonder whether a switch to this uniform pricing would have other welfare reducing effects. Maybe firms would reduce the frequency of flights, they would have worse service, they would put their flights at worse times, cheaper slots and so on. So maybe a bit of caution when interpreting that result.

And the last thing that maybe everybody has on their minds when you read a paper on airlines, most likely a lot of these buyers are frequent flyers. And for them maybe they don't look at competitors' price because they're already locked in into this airline and they like their upgrades and they like going to the lounge and so on. So I actually asked Kevin this morning, they have some data on loyalty and it would be interesting to understand how loyalty affects all of these. It's also interesting to know whether people switch between airlines and how frequently that happens. Yeah. Anyway, that seems something that it's worth considering because it seems likely to affect pricing dynamics significantly. Anyway, I'm running out of time. Thank you.

Speaker 1:

We again open the floor to questions from the audience. Another note, if you do intend to ask a question, please stand up for the videographers.

Tom Kotch:

Maybe I'll start with one. Also need to introduce yourself, so I'm Tom Kotch from the Federal Trade Commission. I was wondering if the new intuition that you develop in this paper, how that might inform our thinking about entry games or in case, DOJ folks are here, mergers in airlines, given you do seem to find again some new intuition about the effects of competition in this space.

Kevin R. Williams:

Should I gather all questions or sequential? That's a deep question that would require a lot of thought on my part. But one thing that we've learned from this research is, firms might not think about their competitors in the ways that we write down. And we discuss this in the other paper and that has pretty big implications on our simulations. If it turns out all airlines don't account for cross price elasticities, where our simulations do, then there's a disagreement with how we study them in their practices. So it's very important. I would have to think about it more in the context of this particular paper though.

Ben Kassner:

Hi, Ben Kassner, FTC. This is a market where you tend to see a lot of intertemporal price discrimination. So you have leisure travelers becoming more price sensitive over time, business travelers becoming less price sensitive. I'm curious how that interacts with this strategic business model differentiation. So the story I have in mind is a route with a mixture of business and leisure travelers. You have one airline with increasing prices over time, one airline with decreasing prices. I'm curious to see whether this model would predict patterns like that?

Kevin R. Williams:

Great question. So that idea of business and leisure and how that changes over time is very much in this paper in a different way, because having two discrete types of consumers, we get to this multi peak thing again and that's problematic for simulations. We do capture it by this time varying price sensitivity though. So it is driving our result and that high willingness to pay arrives at the end and that affects everything firms do. Your question makes me also think about how different firms, airlines have positioned themselves and more specifically dynamic versioning, because airlines can kind of costlessly damage their own product. I have an economy taken but now you need to pay extra for a seat assignment, etcetera. And so we're working on dynamic versioning and I think it's a very interesting question.

Rachel Soloveichik:

Hi, I'm Rachel Soloveichik, from the BEA. I was fascinated by your uniform pricing thing because that's a lot closer to how the federal government negotiates prices for flights of its own employees and wondering if that's how that falls into your model?

Kevin R. Williams:

So we don't consider that here. In separate work we have some discussion on this because we can see everything for one firm. Exactly which tickets are purchased by what types of entities. An important question, can't really tackle it here. We don't observe everything. But I will say broadly for the markets that we study in both papers and probably for the foreseeable future, those sorts of contracts impact international markets much more than what we're studying.

Phil Nelson:

Phil Nelson, former FTC, now Secretary of Economist. A key parameter obviously is the capacity constraint. And my suspicion, but I might be wrong because I don't know the institution as well as you do, is that the capacity, the ability to add additional flights varies significantly across airports, adding a flight to the New Haven airport where I flew is a lot easier than Washington National. So what did you do to think about controlling, assuming that you can add flights more easily at one airport versus at another?

Kevin R. Williams:

Excellent question. We don't do anything here. K is just given and let's study the dynamics given K. But the K question is very important, the influence of airports and slots and actually maybe not New Haven because New Haven has one jet bridge and so maybe you can't actually have that many fights because there's only one bridge. And so it's just an important question that we're not doing here. Relatively how substitutable is demand across airports. I mean, depending on what analysis, sometimes certain airports in major cities are excluded and then the other side says, no, they're all the same. I think we just need more work on this.

Jin Jin Jin:

Yeah, Jin Jin Jin, from University of Maryland. This might be a technical question. To what extent you're take into account the demand substitution between say different time of the flight for the same origin and destination?

Kevin R. Williams:

There is a time of day preference in the demand system. It's one of the X's

Davis Rabel:

Davis Rabel, Federal Trade Commission. You have a setting where your airline knows both P and Q at all times for all of its competitors. Do you have anything to say about the welfare effects of restricting the information that the airline would have about either the quantities of prices that the competitors have?

Kevin R. Williams:

It might not be for all airlines and it might not be for all types of tickets. That's part of our selection criteria. But also what we found in other work and have now told everyone is those things aren't considered at all in pricing decisions. Like loyalty is not considered. The incorporating the P and the Q of competitors, not considered. So I would maybe rephrase that of what's something we're doing in this paper as what if they were? That's the benchmark model.

Speaker 1:

Okay, thanks again to the audience for a great set of questions. For those of you that are here, we'll take a break before our next session at 11:15. For those of you that are online, we'll also be back at 11:15. For those of you here, we do still have coffee and water. For those of you that are online, we do not have that for you, but we look forward to getting back.

Doug Smith:

So, hello again. I guess we're just about ready for our first keynote speech of the conference. So standing in for Dirk Bergman, who was unfortunately unable to attend due to unexpected circumstances, we're very lucky to have Alessandro Bonatti, as a first keynote speaker. Dr. Alessandro Bonatti is the John Norris Maguire Professor of Applied Economics at the MIT Sloan School of Management. His research studies the impact of information technology on firms, online advertising in pricing strategies, as well as on consumer welfare. His most recent work explores the role of data intermediaries with market power, how they collect, mine and monetize information. He's also studied the optimal provision of incentives and research intensive and creative industries and the resolution of conflict inside standard setting organizations. In addition to teaching, Dr. Bonatti serves as editor of the Rand Journal of Economics and associate editor of the American Economic Journal of Microeconomics. Dr. Bonatti will now give a keynote speech on data competition and digital platforms. So please give him a round of warm applause.

Dr. Alessandro Bonatti:

Thank you Doug, and thank you all for being here. Dirk sends his best and his apologies, he has sad but very valid reasons for not being here, so you're stuck with me. On my end it's no less of an honor to be able to talk to you today. The paper that we wanted to show you for this conference is a model of digital platforms and digital commerce. So let me give you some quick motivation.

Okay. So at a 30,000 feet view of this, digital platforms act as gatekeepers of information and managers of competition. They not only sit on vast amounts of data, but they also control which firms, sellers, advertisers, influencers, have access to consumers via the data. And we all know that this has the potential of creating vast amounts of social and consumer surplus by matching buyers with the right sellers, consumers with the right messages. At the same time, this has raised concerns over the potential for extracting surplus from consumers. The idea being that if a platform has market power and somehow transfers that market power to the seller side, that might lead to consumer surplus extraction. This position has been presented in academic work as a possibility. This is a report by Jacques [inaudible 01:39:36] Walters, a few years back for the EU and it has somewhat influenced the Digital Markets Act that went into effect in the EU on Tuesday.

Once we zoom in a little bit closer though, we think that the driver of both the creation and the potential extraction of surplus is really the personalization of sponsored content. Sponsored content, you can think of Google Ads, you can think of sponsored listings on Amazon and a bunch of other platform is a growing fraction of the revenue of retail platforms. Think about Amazon's advertising revenue to name one. It is a hundred percent of the revenue of retail platforms in China like Taobao, that does not charge fees. And it is obviously a hundred percent of the revenue of pure advertising platforms that are not selling anything. They're just selling ads. And when we talk about sponsored content, we include targeted advertising.

So it's enough for me to go on social media and be shown my favorite coffee or my favorite executive education program that Warton wants to enroll me in, or on other sites. But I do want you, as we go to the model, then to keep in mind that these are Google advertising network, Meta advertising network, Microsoft, that this doesn't want to be a talk about Amazon, it wants to be a talk about sponsored messages on a variety of platforms, including display networks. Those that place ads on third party websites that are not Google sites, not Facebook sites, they're just placing ads on their behalf.

So when we think about it in this way, there are really three ubiquitous features of digital markets and digital commerce. One, the platform has an informational advantage that it leverages and monetizes. Two, offers are personalized, but as far as we know, the evidence of pure personalized pricing is very scarce. What does tend to happen a lot more is product steering. Julie and I don't pay the different prices for the same good. We are shown different things. And third, however personalized content is sold and it's ranked, it's going to be sold through sort of an auction mechanism. And this want's to be the last bit of institutional detail that I want to discuss with you, which is what are these auction based mechanisms?

When I started teaching game theory to MBAs maybe 10 years ago, I was teaching them the generalized second price auction and how you would manually bid for keywords on Google. Hardly anybody does that anymore. Okay. Ads are largely sold through automated campaigns, auto-bidding, managed campaigns. They're all synonyms for our purposes, which really mean, give a budget to Google or Facebook and they're going to optimize the spending for you imputing fractions of the budget to each ad that they serve. Learning through their algorithm, which one is the best consumer profile that you are going to be exposed to without you having to manually adjust your bid keyboard string by keyboard string or website by website. Okay. Does this matter? We'll see. Okay.

So with this context in mind, we want to answer a few pretty broad economic questions. One is, how does the precision of the platform's data affect the creation and distribution of surplus, not only on the platform itself, but in the rest of the ecosystem? Do platforms with market power have incentives to provide high quality matches? Intuitively, yes. Do they, or can they transfer this market power downstream, create these monopoly positions that the primary report was worried about? And to be clear, this is the FTC, as far as I know, there is no evidence of these positions having been created, but it's definitely a plausible concern, at least for theory.

And then fourth, what is the role of the platform's revenue model? Does it matter that they're charging fees or manage campaigns or manual bidding campaigns? So there are other questions like the ones in gray, like self-referencing that we could answer in this framework. We don't as of now. Instead what we do is provide the model with four key features. One is, there is an urgency all over the place on the buy, on the demand, and on the supply side. There is personalization of content, there is products steering, but there is no price discrimination at the personal level.

And fourth, and maybe the newest element of the model is there are parallel sales channels. The platform is not the only place where firms, I'm going to call them firms, can be sellers, can be advertisers, influencers, are going to find their consumers. So those are our four key features.

In the interest of time, let me not describe any papers in particular. There is a growing literature on multiple sales channels. There is older literature on information gatekeepers, beginning with the seminal work of BAE and Morgan. There is a grayed out literature on self preference in problems including Z Yang's paper from last year. And there is lot more references in your paper. We're about to finalize the paper. So if I miss something in the paper, not on the slide, of course I'm very happy to hear about it. And also there's no discussion, so I don't know how you want to

Dr. Alessandro Bonatti:

Around this thing, but if you have clarifying questions, just please ask them out. It's, I'm going to go in through the entire model. Okay, so while I've given you I think a bunch of institutional detail and talked about competition and the focus is on how the platform is going to manage competition and information, I would like you to forget a little bit about that and want to take you back to Grad IO or first year micro course and do a simple single seller quality pricing example with a twist that's going to capture what we do in the model. I'm going to talk at length about the example and then bring competition back and show how a platform can limit competition in this model.

So this is kind of the engine of the model in the simplest form, and then I'll add all the bells and whistles. Okay. So, go back to your second degree price discrimination or quality pricing. There's a single seller, like in Mussa-Rosen faces a unit mass of consumers. Consumers having two types, the low type and the high type. They have a given distribution, it's a binary distribution. Preferences are linear. The consumer of type theta has valuation theta times Q for product of quality queue goods can be produced by the seller quality queue with a quadratic cost. So it's textbook example.

Now, one minus Lambda of these consumers independently of theta are going to buy directly from the seller. So that's literally the textbook example. A fraction Lambda exogenously visit a single platform that can run ads.

Now the platform has a lot of data. We're not going to model how they acquired it and is going to offer a managed campaign to the seller. What is a managed campaign to the seller, is the right to personalize offers at the consumer type level. So how does it work? The platform charges an upfront fee, T. call it a campaign budget if you want. It asks the seller, Okay, what product qualities and what prices would you like to charge to each of the two consumers type the low and the high type and then promises the seller that whenever consumer type theta shows up, the platform is going to show the relevant product, the one intended for type theta to that consumer or equivalently the product that maximizes the consumer's profit given that type and given the knowledge of the type that the platform has. In parallel and in fact simultaneously the seller is also going to post a menu of products on their own website or SOAR or whatever channel they have.

And it's going to have to be a menu because off the platform, the seller does not know the consumer's type, so they're going to have to screen them. So think of it as your textbook model of the platform and personalized offers that the consumer will not be able to choose among, that the consumer will not need to self-select between on the platform. And you could think of it as a constraint, but now let's leave the platform free to show the consumer products that are special only to the platform and they are not available on the seller's own store that could be easily introduced as a constraint.

So schematically, the seller's going to have their own site. They post a menu of products, it's really two qualities and two prices with two types on their own website. And then the seller can also upload this customized offers to the platform. The platform does this auto-bidding thing, but there's only one bidder. So it really, it's going to select one offer to show to each consumer type that visits the platform. That is the one that we call the online consumer top right. The offline or off platform consumer, bottom right goes directly to the firm's website.

And the only connection between these two worlds is that if you were a platform user after having seen your offer, you can also check the firm's own store. say, "Oh, I saw the brand of coffee, it exists, great. Let me check what they have on their own website." If you are not on the platform then you just live off the platform. So the dash line is really the only connection.

And the question is, how is this going to affect potentially these four products that the firm is going to offer? There are two types and there are two channels. So in principle, they could offer for a product for the low and the high type on the platform and maybe the same products or maybe different ones off of it for each of these two types.

Can I... yeah, can I go back? Okay. If something is not clear, please ask now. Otherwise let's just think it through. Okay.

The platform could open up, its... Sorry, the seller could open up its graduate textbook and figure out that they have to solve this simple two type exercise on their own website and find the Mussa-Rosen solution. And the high type is going to get a high quality and a high price. The low type is going to get a distorted quality downward and a price that extracts all its surplus, efficiency at the top, no end at the bottom, and that they could do on their own website. And then on the platform, the platform knows the type, so it really trade occurs under symmetric information. So they might want to offer the socially efficient product to each consumer on the platform and potentially even extract all the consumer surplus because they know the type via the managed campaign. So that would be ideally what the seller would want to do.

Of course that doesn't work because if you're charging the online consumers too much, they will just check what you have in store on your own menu. And if they get a positive utility out of that, they would, technical term they would showroom, they would use the platform just as a listing and then go and buy from your own website.

So again, ideally the platform would offer the socially efficient quality that is in this quadratic model is just Q of theta equals theta and charge the consumer's full willingness to pay. And they could do that if they didn't have demand on their own parallel channel. But they can't do that if they are serving people on their own channel too, which means that the showrooming constraint is going to be very simple. It's going to be u of theta, the information rent that the consumer obtains on the platform needs to be greater than or equal to u hat... Hats are for off whatever this consumer could get by showrooming and choosing their favorite product from the publicly available menu. Okay? That's in addition of course to the fact that the public menu needs to induce self selection so it needs to be IC and IR for the consumers that visit it. Okay?So the bottom line is the platform via their own data in this managed campaign enables trade under symmetric information. And so some form of product steering and price discrimination, but the extraction potential is limited by the publicly available offers that the seller also has. Okay, so what is the solution? The expression is long, but there's really only one part of it that I care to show you. What does the seller do? The seller offers the efficient qualities on the platform. Showrooming binds. There's no reason to make consumers strictly happier to be on one channel versus the other.

It offers the efficient quality to the high type of the platform. That's no distortion at the top, that's classic. It distorts the quality to the low type on the platform. That is also a classic result. But you will see that there's a lambda over one minus lambda term and that has a minus sign in front of it. That means the larger the platform, the customer base of the platform is the larger the distortion in the quality provided to the low type off of it. And if you think about it for one more second, distorting the quality of the low type of platform is going to reduce the information rent of the high type of the platform and it's going to relax the sharumi constraint. Okay.

So before I show you a picture, let me just point out that the line at the bottom which says there's always an opportunity cost of putting a rich menu on your website. This is tier [inaudible 01:53:18] Rosen. It's the fact that I have these option for the low types. I need to give a rent to the high types. That's always there. There is an additional opportunity cost in this model, which is any rent I give to anybody off my platform, I also need to provide it by offering lower prices on the platform else they wouldn't buy there and my ad money is wasted. Okay, let me put it like that.

So let me spend maybe a couple of minutes talking through the picture, which says everything that I wanted to say. You have qualities on the left, prices on the right. The independent variable is Lambda, the size of the platform, the fraction of consumers that you can call them double homing consumers, dual homing because they visit both channels. So in blue you have what happens to the high type. The high type always gets the best product. Q equals theta. That's the efficient quality regardless of where they live. And they always pay the same price. Well, it can't be otherwise. If you were charging different prices for the same product, the consumer would just pick one over the other.

The low type on the platform gets the red line on the left, which is the efficient product for her type theta L, and pays a price which is theta L square, which means that consumer makes no rent on the platform. Why does it make no rent on the platform? Well, because that consumer also makes no rent off of it because they're the bottom type. However, off the platform, they would not have access to the same product or put differently. The red line product is not available on the firm's own menu. There is a worse product that's available and it's getting worse and worse, the bigger and bigger the platform gets.

It's offered of course at a lower price, but also that lower price extracts all of the low types surplus. So this type makes no ends regardless of where they go. So if you want me to tell you a story for this, it would be as if you go on booking.com or some travel website and you say, Oh, here's a special offer for you. Get an upgrade to your hotel room at this sort of advantageous price. If you were to buy from my own website, you would just find a worse room at a lower price and you weekly prefer the upgrade that's personalized just for you.

How about the comparative statics here? Well, the more consumers are on the platform, the more the platform cares about charging a high blue price. The top line in the right hand graph and the cost of charging a higher price to the high types, both on and off, is that you need to reduce their outside options. And you do that by offering a lower and lower quality to the low type on your own site. Okay?So the classic distortion that's always there for price discrimination purposes in a quality pricing model becomes much worse here and becomes worse and worse the more you care about extracting the high type surplus on the platform where you actually get to make them a personalized offer. It's basically take this or leave this and you get to charge a higher price.

So this is what would happen, there's only two types. There is a, there's a single seller. What can the platform get? Well, the platform cannot quite extract all of the profits of the seller. The seller has an outside option. The outside option is, okay, I'm not going to advertise. I have one menus lambda consumers in my store anyway. I can solve my [inaudible 01:56:41] Rosen, my two type example.

I cannot distort that menu further because I only have one channel and I can always score those profits. So this means that the platform T star, the optimal fee for the platform, which of course is going to make this single seller indifferent between participating and not has to be strictly lower than the profits that seller is going to make on the platform. It has to be strictly lower than that because in order to advertise on that platform, the seller optimally distorts their pricing in their own store away from what the monopoly solution would be.

So there's a little bit of compensation there. So if I were in the marketing department of one of these platforms, I would tell them the campaign delivers a positive ROI because the revenue you make on the platform is strictly greater than T star. Now of course I'm pitching this as a marketing pitch because at the end of the day, the seller is held to their outside option.

But if you really split the accounting in this way, then what I said was strictly speaking be true. Yeah. Okay. So what have we seen? We have seen that the platform uses its information to enable efficient trade. There are no distortions. Everybody gets the efficient quality on the platform. They extract everything that they can extract. The result is higher prices for the high type off of it, and in fact lower rent levels for every type regardless of where they live. And then they can extract whatever they have induced and created for the seller through this fixed fee upfront. So this is sensing how the managed campaign differs from a bidding. The seller is now not bidding for this service, it's a mechanism, it's a take it or leave it offer that holds the seller at their outside option.

So, in the remaining 14 minutes and 35 seconds, I don't want to solve the whole model for you. I want to show you, in which sense is the example too easy and how does it generalize however, to a more complicated world with competition? The example is too easy. Yeah, there are only two types that we can relax, but more importantly, there was only one seller. And also I made a big deal out of the platform having a large informational advantage, but the consumer really didn't need to use the platform to figure out anything about their preferences. And this is an important dimension of value creation, of figuring out what are things that we like.

The example is also a little too hard in the sense that I could have told you a one quality story where the platform enables personalized pricing and some of the economics that we just discussed would've gone through in that model. But then many of you, especially the more applied colleagues, would've told me that there is no price personalization online. And so then I would've had to tell you that, oh, then we have a product steering story that's kind of equivalent. And so I decided just to put up the product steering story first as a maintained and of factual description of the world out there.

So if we have time to talk about extensions, and sometimes you do want to go to the unit demand and unit quality example because that of course allows you to go farther in some dimension. So the name of the game for the rest of the talk is to show how and under which conditions can our monopolies platform recreate the single seller world you just saw, despite the presence of multiple potentially competing sellers.

So I want to get to a place where the equilibrium basically looks the same or I want to find conditions under which the equilibrium looks the same in a more complex and more realistic setting. So we had one set, one Mussa-Rosen seller. Now we're going to have competing Mussa-Rosen sellers. So that means that the consumer's type is still theta, but theta is now a vector. And the end, the J component of this vector is the margin of valuation that this consumer has for a quality produced by firm J. Okay? So it's a differentiated products extension of Mussa-Rosen. Firms are symmetric, etc. They all have the same cost independently of theta. Lambda consumers are on the platform. One minus Lambda are off.

The consumers valuations are independent across brands and consumers. I don't think that's super important, but it will maintain that. The platform knows the entire vector of taste of each consumer and consumers are betrayally, uninformed about their preferences. They only observe a signal S that leads to an updated belief, MJ about your valuation of quality for firm J. If you want to think M is almost equal to theta or almost surely equal to theta, you can, but this informational advantage that the platform has over the consumer will play a role. And surely F is a mean preserving spread of G because, because M's distributed according to G, are beliefs about theta.

How does a managed campaign work when you have more than one advertiser? So formally speaking, the platform charges posts a fixed fee and says everybody wants to participate in this campaign, pays T and tells me I'm speaking as the platform now... Tells me which products and prices they want to show to which consumer theta, just like we saw in the example.

The mechanism now also specifies which brand and is going to get the advertising slot for each individual consumer. And what product from that brand is going to be shown in the ad to that consumer. And I'm just going to put it up in the interest of time, the optimal mechanism for the platform, the one that maximizes their equilibrium revenue is going to be to do the efficient thing, which is to commit to showing the product from all the... There's a continual products for each firm and there are capital J firms to show the point product that maximizes total surplus for that consumer.

So this is going to be the consumer's favorite brand and the product intended for that consumer by that brand. But that's in equilibrium. If you want to state it as a mechanism, it's I'm going to show the product that maximizes social surplus

And you can sort of see where I'm going. The platform is committing to generating as much of the pie as possible and then they're going to try and extract it with T and it will work.Now, off the platform, they're competing brands. So how does the model work? The model off the platform works like the DIA model. So consumers, they have these types really, M, their beliefs, but that's all the information that they have. They face search costs and the first search is free. So everybody searches in Diamond. We think these costs are substantial, but there's the diamond paradox. You can think of them as being as small as you would like. It's still going to work off the platforms. The seller, remember they don't know the consumer's type. They only have the distribution of consumer's taste. So they're going to have to post a menu to screen the consumer's type, or in this case the consumer's beliefs.

Okay, this is not an inspection good. It maintained... Assumption is that the platform has this strong informational advantage and they are going to, they're the only ones who can reveal the consumer's true taste. But then again, I told you that M can be arbitrarily close to theta, so I wouldn't worry about that just yet. The timing, the mechanism is announced, then everything else happens simultaneously, posted menus, uploaded ads, participation decision, and then of course a consumer is realized, the ad is shown, the consumer decides what to do.

So the only result that I want to talk through you at length in this richer model is what happens in equilibrium to a consumer's search and behavior patterns.

So, off the platform, it's by design, it's the diamond model. So the consumer expects symmetric menus to be posted by all firms, but they still need to search to figure out what the prices are. They have an arbitrarily small but positive search cost. So, they know that their rent function, informational rents are going to be increasing in their type. So really they do the natural thing, which is, oh, I have a firm that I think provides the highest value to me, let me search for their menu and pick out whatever product from their menu is maximizing my utility. So it's a self selection story.

But either way, they're only going to visit the firm that they think is best. What happens on the platform? Well, on the platform, the consumer is going to... Let me just do that. On the platform, the consumer at the top is going to see an ad and imagine that she sees an ad by firm two.

Now the consumer knows the platform has a small but relevant informational advantage. So the consumer deduces, that firm two is in fact the one that maximizes her utility and she learns theta two, the value for that product. At this point, the consumer can buy the personalized offer that she just saw or she can search. But if she searches, she has now learned that brand two is the one that maximizes her value. So she's going to go to firm two's website and compare and show room if the prices are not right. Now Counter factually, if the same consumer lived off the platform and had beliefs M. Maybe those beliefs were wrong. Maybe the consumer started out thinking that firm one was the one that maximized her utility, in which case she would've gone to from one's menu. And that because she lives offline and she has search costs, that's the only one she would've seen.

So to put it differently in equilibrium, our consumer is only... Online consumer, is only going to contemplate two products. The one that she's advertised and her choice from that same brand public menu. And because the platform has a small advantage and there's a small search cost, the online consumer will never see or visit any other brand's menu that is not advertised on the platform's website.

So to put it differently, part of the concern that or other colleagues have had was that the page that all the other relevant competing search results will be hidden on page two of the search results. And that was the joke on the slide. But to say it more formally, if the platform is as strict arbitrarily small informational advantage, the consumer is only going to compare the displayed seller's offers. And this is going to be quite powerful because it's now going to tell us that as a brand, if you now think about foregoing this managed campaign and turning down the platform's offer, well you are not going to see any of the Lambda online consumers.

What will the Lambda consumers see? They will see the second best firm. They will think that they will not be able to detect the deviation. They will think that the second best firm is in fact the best because the platform knows better and they will not come back to your website if you haven't advertised. So, this will be the combination of small search cost or in a small informational advantage, which is what I just show you, turns the J capital J firm problem into the continuous types version of the example that we just discussed. To the point that I could have started with maybe less micro founded, but or more behavioral, but completely equivalent interpretation. I could have said every brand has a woman Islam that divided by capital J loyal consumers, they're going to buy from your website no matter what.

Lambda consumers, they're not shoppers, they're not thinking about buying anything. They don't know that brands are out there. If they see an ad, they learn their value for that product, they understand that given brand of coffee exists and they're like, okay, let me go and check out this brand's website and see what else they have in store. And this more consideration sets or awareness or activation interpretation that I did not start with would deliver of course exactly what we get in the equilibrium of our search and information friction model.

So I guess you can imagine where... Because we saw it in the first part of the talk where all of this is going, this is going to a world not to an equilibrium, then I can describe it to you in words where each online consumer is going to find an equilibrium. Everybody participates is going to find their truly favorite brand. They're going to buy the efficient quality from that brand because trade occurs under symmetric information. They're going to pay possibly a higher price that they would've paid if they show roomed, but they're also having access to a better product, thanks to the fact that the platform's information removes the need to distort quality downwards.

They're indifferent between the two, but they're indifferent between something that's great at a high price and something that's a little less great at a lower price. If they showroom, what happens to the off platform consumers? Well, they suffer in the sense that qualities are going to be distorted down just like QL hat was distorted down in the example. And why is that? Well, to reduce the rents that are provided off the platform and to raise the prices that are charged with the advertised products on the platform.

So the setting that I described is one of endogenously local monopolies. Thanks to these two factors, search cost and information friction, I never... Let me say it differently for the IO and competition colleagues. We never had to solve for an equilibrium of a competitive pricing or competitive non-linear pricing game because in this equilibrium of this mechanism, which does maximize the platform's revenue, you are only competing with your own offline store. And then the size of the platform, the size of these consumer populations tells you how much you're going to compromise your offline store. Let me put it like this, in order to charge higher prices on the platform. And there's an interior resolution to this trade off that goes in the same direction that the first two pictures went. The bigger the Lambda, the less you care about the direct sales channel, the more you distort trade, you might even shut down a bunch of types from trade in that sector.

Okay? So what we are setting out here is a model where a digital platform monetizes its superior information by auctioning off access to the consumer's attention. In my other interpretation in a very specific way through these managed campaigns. what we do in the rest of the paper is try to convince the readers and you all that our two ingredients are in fact necessary and sufficient.

So the world looks a little different if the consumer knows her type perfectly and can't be "fooled" by a message sent by the platform. And it looks radically different if the offline menus are publicly posted. That which you might call organic search links on Google. And the world looks yet again different if you limit the platform's ability to condition the message on the consumer type. For example, you limit it by only allowing ads that based on the consumer's rank of brands and not the precise quantitative valuation for each product.

Now I think something that I have 20 seconds to talk to very quickly about is something that I don't have any results over, but is this idea of a trade off between privacy and competition. So what we saw in this paper is that competition is entirely restricted in this mechanism, but the mechanism is quite private after all. There's only one brand that accesses the consumer at one time. They never get to theta, they don't even know how much they spent from their budget to access their consumer and they never learn anything about this consumer elsewhere.

So it is a pretty private mechanism. Is this trade off always there? Well, not really. We're working on a follow up with Nick Wu, who's a third year student at Yale in which we compare this managed campaign to the arguably less private auction mechanism where each brand is conditioning their behavior exactly on the consumer's type. And preliminary results show that actually the managed campaign might be an improvement on both fronts. So I know this is a hard question and one that I'm interested in personally, and there's follow up work to this paper that possibly can speak to that. Thank you very much.

Tom Kotch:

So thank you very much. It's lunchtime. So lunch for those people that are here in person has been provided to us by the Yales Tobin Center. We asked that if you were here in person to not leave the building. It is a beautiful day outside. That's a trap. If you leave the building and go past security, you'll have to go back past security. So again, if you're here in person, just have lunch in the local area around where we are now. If you are watching us online, we're not providing lunch for you, but you are free to go outside. Either way, we will start again at one o'clock. So we look forward to seeing everyone in person or online at one o'clock. Thank you.

Come back everyone. We are now prepared to start our next session,

But before we do that, I just had a couple of notes on by behalf of myself, Tom Kotch and my co-organizer Will Violet. We wanted to reiterate the thanks that Mike gave to the BE staffers who help put us together. And in particular we wanted to reemphasize our appreciation for Stephanie Aaron. For those of you that have been involved more closely with the participation in the conference, you have received a lot of organized, thoughtful, and well put together emails from her. Know that Will and I are incapable of doing anything remotely like that and that without her effort and participation and thoughtfulness, there will be many, many more seams present in this conference. Another note. So Stephanie has a job that was previously done by Alex Arama, who left the commission last year to find another job. He left incredibly detailed notes that enabled us and Stephanie to do the good job that she's done.

So Alex, if you're out there on the internet, thank you for making our lives easier. And then the final note I wanted to mention was the events staff that have helped put this on. So, this looks a lot more complicated than most conferences that I go to in part because it really is with all of the [inaudible 02:16:22] and this tables and the sands and the technology. Our three main points of contact have been Arisa Henderson, Jacqueline [inaudible 02:16:29], and Bruce Jennings. A thing to note, two things to note. This is the first in person to my understanding meeting that the FTCs held since Tom Hanks got COVID. So it's been several years.

So everyone has just kind of tried to refine themselves in doing this and they've done an incredible job managing this and helping us helping do this. Kind of relatedly for those of you who are maybe in the privacy space, they also held an organized privacy con, which was just a couple of days ago. And I was talking to somebody at the commission level, I said, "Oh, our conference for Microeconomics is on Thursday and Friday." And she said, "Tom, privacy cons Tuesday. What are you doing to those poor people?" We are really kind of testing the limits and they are doing an incredible job of helping us put on this conference and extraordinary circumstances. So I just want to make sure that they hear from us. Thank you.

So if you've been enjoying your time here, you think things are well run, don't thank Will, and I thank the staff and the people that we've just listed.

Now back to research. So this is a session with paper selected by Will and I. The first paper we presented by Michael Richards. It'll be discussed by Lemore daphne. The Paper is entitled Outside Equity and Healthcare Firm Behavior. A note is that Lemore will be joining us via Zoom, so half we'll be paying attention to the screens when she discusses. Thank you.

Michael R. Richards:

Okay. Thank you all for being here and thank you especially for the organizers, the staff, everyone that put this together. This is joint work with my colleagues, Hygen Lynn, Beth Munich, Chris Waley, and [inaudible 02:18:25] . And we are happy to present some early research, our early results from a relatively new paper for us. Just in case we don't get through everything, there's any interruptions, I know we had warnings about fire drills and other things. This is going to be a paper about Ambulatory Surgery Centers or ASCs and in particular their behavior under different financial investment contexts.

If you don't know what ASC is, don't worry, we're going to spend a lot of time talking about that. But we're going to find that private equity investments in particular are mostly going to engage in financial engineering of these firms as opposed to trying to change their treatment styles or behaviors. At least among these standalone ASC firms. In particular, they're going to start charging more per case and they're actually going to crowd in physician equity investors as opposed to crowding them out.

But when we look at the private equity acquisition of a large prominent ASC chain, things can be a little bit different. And so we can't assume that that private equity investment will play out in exactly the same way in all contexts. And then interestingly enough, as this chain is approaching its IPO, its initial public offering, there does seem to be some behavior change that could be consistent with revenue maximization for the valuation. Okay.

So something that has occurred over the last 20 to 30 years is a shift from a lot of companies throughout the US economy raising capital from private markets as opposed to public markets. And these private

Michael R. Richards:

... investors or entities come in a mix of things. It could be venture capital, it could be real estate investment trusts, and so on. But one of the major players that looms large is private equity. So they've been involved in a lot of different deals, and they typically have a lot of capital that they're able to deploy in these ways.

Now, the views on private equity's involvement in different companies and sectors of the economy are mixed, and the empirical results seem to be that way as well. So in some findings suggest that maybe they improve company performance and enhance valuations, but others find that at least for key stakeholders, they may be worse off under private equity control.

One place where private equity has been having a larger influence over time, and also creating a lot of controversy is healthcare. And so over the last roughly decade, about $800 billion of private equity funding has flowed into different US healthcare companies. In 2018 alone, it was about a $100 billion, so this is something that is relatively aggressive. And at least among industry insiders, one explanation is it seems to be that the investments that they make in the healthcare space tend to outperform those that are done in other industries or sectors of the economy.

I may have gone too far. Can I go back? Yes. So all that is to say that this is something that is growing, but something that we don't have a lot of insight about and something that a variety of stakeholders, regulatory entities, and so on, are trying to understand. What are the implications of all of this money from these particular firms coming into the sector?

And one of the reasons that healthcare in particular can be a source of controversy and private equity's involvement, is there's a concern that will private equity investors, because of their business objectives, they're relatively short time horizons for their financial endpoints, for example, selling a company within five or six years, will this lead to behavior changes among healthcare providers that are potentially at odds with what's the best interest of the patient? So will they engage in different types of rent seeking behavior? Or will they distort a provider agency away from what might be optimal from patients?

There is an emerging literature in this space. A lot of it kind of focused on hospitals, nursing homes, and physician practices. But some other key industries are less investigated. And so the one that we are most interested in today is, as I said before, the ambulatory surgery centers or ASCs. And actually private equity has a pretty long history of investing in these firms, but there's not an empirical literature to really come alongside it.

In fact, we're only aware of one very, very recent and new study in this space. They used a sub sample of Medicare data over a relatively short analytic horizon, and focused on the quality of care and the health outcomes that happened after private equity investment, but didn't really find much of effects. So we're going to take a little bit of a different approach. And so we're going to leverage some all payer data over a relatively long time horizon, so that we can look at both the investment and divestment endpoints. And then been taken advantage of this particular context where we can examine an IPO. And we're going to look for changes in payer mix, case mix, treatment intensity, list prices, and then physician ownership of these same firms.

It's my understanding that not everyone studies healthcare. So now we're going to do a lot of institutional details. The first thing to be aware of is that medical care is very much moving to the outpatient side. Originally, with this bullet point, I was going to make a Billy Joel reference about moving out, but then I realized that this is being recorded and could live forever on the web, and so the expected embarrassment costs were just too high, so we're going to have to move on.

Even on the hospital industry, about half of their revenues are coming from the outpatient side, alongside the inpatient side. And where we're seeing these shifts really most pronounced is in surgical care. So a lot of procedural surgical care is now able to be safely and efficiently taken care of in the outpatient side.

But there's essentially two rival industries that compete for a lot of these cases, not all of them, but a lot of them. And the one is the ambulatory surgery centers or ASCs. And then the other is the hospital outpatient departments, or sometimes called HOPDs. Now, it's not always easy to know one from the other, and so if you've been driving down the road and seen something big, flashy, and impressive, that's probably a hospital outpatient department. If you've seen a relatively small mom-and-pop shop, probably, situated in a dying strip mall, probably an ambulatory surgery center. So kind of think of those two extremes to keep in mind what this market looks like.

ASCs tend to be small and for-profit, and so on average maybe two to four operating rooms or theaters per facility. They tend to go to the urban areas, and they've grown quite a bit. So in the early '90s, there was about a 1,000 of these certified by Medicare across the country, but now we're seeing over 5,000 of them spread throughout the country. It's estimated to be about a $30 billion industry, so it's a non-trivial player in the healthcare space. And even Medicare, specifically, spends about $5 billion in care in ASCs.

The value proposition for ambulatory surgery centers is they typically are more convenient for consumers and they offer lower cost care. And so it makes it an attractive option. So not surprisingly, in contested markets, where you see more of an ASC presence or entry, HOPDs tend to fare worse, so they lose business and suffer more financially. Helpfully and usefully, for our purposes, private equity as well as other investors have been engaged in the ASC industry for quite some time.

Okay, so this sets up a relatively straightforward research question is, do these outside equity investors, private equity in particular, kind of change the behavior of ambulatory surgery centers? And does the context matter? Does it matter if it's more of a standalone firm versus this horizontally integrated chain? And does the arrival of an IPO change the behavior?

Fortunately, as you'll see in a moment, we're going to do all this through a reduced form estimation. So you're not going to have to check any math during this presentation, which is quite nice after lunch, I think it aids in digestion. So we're going to bring in data from a few different places. So the first of which is going to be from a FOIA request to CMS, and this is going to give us very granular, detailed ownership information on ambulatory surgery centers. So for all these ASCs across the country and over a long time horizon, we're effectively going to be able to reconstruct their historical ownership profile.

We're then going to use a variety of different resources to try and identify specific owners that are private equity firms, so that we can kind of tie those to these ambulatory surgery centers in the state of Florida where we have the data. And then we're going to leverage a universe of outpatient, all-payer encounter data over about a 15 year span. Now, we are having to trade-off geographic scope for detailed granularity. So we're going to be focused on the state of Florida, which means that we may not be able to generalize completely across the country, but we do think that the quality of the data is very helpful for what we want to do.

And for those of you less experienced or familiar with healthcare markets, Florida tends to be a big deal. So, a lot of companies have a lot of strategies that play out in Florida. So all the major health insurers are present in the state. The academic medical centers are medium-sized FSHD rather than dominant entities in their market. And even when new companies are coming online, they'll usually start out in Florida and see how things go, even when companies are retreating. And so they're pulling back products from the market, they'll usually let the Florida one stay. So while it may be one state, it's a particularly interesting one, and one where there's a lot of fierce competition among healthcare firms and industries.

So our empirical work is going to effectively take part in two phases. So the first of which is we're going to focus on the investment and divestment endpoints for these standalone or singular ASC entities. And we're going to do that in a stacked difference-in-differences set up. So as many of you are aware, at least when it comes to the reduced form analyses, the difference-in-differences approach is rapidly changing.

And may be even, that literature may have advanced just while we were having lunch. So we're going to have to be aware of that, and we're going to have to make some empirical modifications to deal with it. But then when we move to the wholesale private equity acquisition of an ASC chain as well as an eventual IPO, we're going to be in a much more standard diff-in-diff setup.

So, when we're looking at the investment endpoint, we're going to focus on incumbent firms that have been in the Florida markets for a long time, and we're going to focus specifically on these ambulatory surgery centers that are going to go under private equity ownership for at least six years. And you can see on the histogram, typically, it's about six to 12 years that they're under private equity ownership. And we're going to compare their behavior and in particularly changes in their behavior to other firms that never have any direct or indirect involvement with private equity, and are also what we're going to be calling out-of-market. So they essentially are serving patients outside of the geographic catchment area that our private equity ASC firms are engaged with.

So all of our data are going to be at the ASC by quarter year level, and we're going to, as I mentioned before, do this stack difference-in-differences approach, which for the treatment group is relatively straightforward. We're going to zero in on the time that private equity acquisition or investment takes place. We're going to look two and a half years before that, and we're going to look up to six years out from that, and kind of line up everybody together.

But then that leaves open, well, what do you do with your control group, since they don't have that same centering point? And so what we do is randomly assign them a placebo or anchor time point based on the acquisition timing that we observe in the treatment group. And then we just take their two and a half up to six years and put it all together, and then it's just implementing a straight difference-in-differences from there.

This is just a quick demonstration or rundown of what the treatment versus control group looks like. In terms of those that will come under private equity investment, they tended to be more productive at baseline and also charging more for the services that they provided. Interestingly enough, the payer mix breakdown among the two most prominent payers in the outpatient market space of privately insured versus traditional Medicare is effectively inverse across the two. So where the firms or the ASCs that the private equity funders seem to target are much, much more tilted toward the privately insured market compared to Medicare.

About half of them have physician owners at baseline, and just one other kind of institutional feature. The vast majority of ASCs are actually physician owned and operated. And so that's still true today, about 70% of ASCs across the market at the moment our physician owned. In the early years of data that we're working with, it's probably more in the 80 to 90% range. So that's the common approach. And for those with physician ownership, it's about two physicians at baseline.

So here now we're going to look at a whole bunch of pictures for our empirical results. So the first thing is, in terms of the actual case volume or productivity of these firms, it's not really changing with private equity. So there doesn't seem to be any behavior that would suggest that they're trying to either bring in more cases or dial up the services. Now they are charging a lot more for the cases that they perform. We need to, again, in the acknowledgement that not everyone does healthcare, I think we need to take a little side conversation of what exactly charges mean.

So for those of you who don't know healthcare, the short of it is most things are ridiculous, pricing being one of them. So we have this concept of charges that come from charge masters, which are probably best thought of his list prices, but they don't necessarily reflect transaction prices. And so, unfortunately, in the data, we don't observe transaction prices, so we can't get there. But in terms of just an institutional detail, we do know from negotiation between healthcare providers and in particular the key intermediaries of insurers, charges do have influence.

And in fact, they're often a point of kind of contract negotiation between these entities. And a lot of the contracts will look like a percent of charges. So they might tell the healthcare providers, "We'll give you 60% of whatever your charges or list prices are," because otherwise it's really, really costly and time consuming to negotiate over every particular service that they may provide. So while we don't observe transaction prices, we could imagine that this increase in charges is actually going to translate to some higher prices paid.

Interestingly enough, while they're not doing more cases and they're charging a lot more, up to 50% more for the cases they do do, they're not doing a greater intensity or more aggressive care. So the actual number of procedures per case is, if anything, falling. And then when we create this kind of complexity measure or complexity index, we see that's largely unchanged over this time period. So they're not doing more, they're not doing a higher complexity care, but they are increasing their list prices quite a bit.

And we also see this increase in list prices across all the different payer markets, which is helpful for just kind of re-ensuring this, that this is something probably more consistent with charge master behavior as opposed to some sort of strategic thing with certain payers. The other thing that I just mentioned here was actually a recent paper from Sing and colleagues focused on physician practices, was able to observe both charges and transaction prices. And they show what you would expect, if charges go up by about 20%, and actual transaction prices go up by about 10%, which is typically about what the cost to charge , or pricing to charges looks like.

One interesting thing to note, when we go to their payer mix, specifically, we see that the privately insured share of their payer mix drops considerably. So they're down about 20% soon after the private equity investment is made. So this is at least consistent with perhaps some private insurers kind of balking at negotiating with these firms going forward, especially if they're wanting higher prices for services.

Another thing to note is while we can't piece apart the different types of private insurers, one thing that we want to keep in mind is, we often talk about it as a very homogeneous group, but in fact, it's probably the most heterogeneous group among insurers. And so these firms are going to have a wide variety of privately insured contracts. Some of them, they're going to value more than others. And so it may not just be the private insurers walking away from the negotiating table. There may be certain contracts that these firms are willing to give up.

The thing that is particularly interesting to us, because my colleagues and I have a lot of interest in the physician equity behavior in the ASC industry, is when these infusion of private capital comes into play, it's not actually crowding out physician investors or physician equity owners. If anything, it's crowding them in. And so both on the extensive margin, more of these firms actually have physician equity owners, but also the number of physician equity owners grows quite a bit. It's up 300% about a year and a half after the private equity investment takes place.

So then on the divestment side, we're going to do a very analogous approach. We're just moving from the investment financial time point to the divestment end point, but do the exact same kind of stack difference-in-differences look. In the interest of time, I'm not going to go through all of the results. Most of it suggest that essentially the behavior changes that the ASCs make when private equity invests largely continue.

The one thing that is different is the physician investments or physician equity stakes, and in particular, at the time that the private equity cashes out or sells their stake in the ASC, the physicians do likewise. And so there's both on the extensive margin and intensive margin, a big drop off in physician equity investments. So the data seemed to suggest that they actually coordinate their divestment decisions. So while the private equity helps recruit in these physicians, they then kind of coordinate and align their divestment time points as well.

Also, worth noting, the physicians that become equity investors in these firms are not actually new physicians. So they were already doing cases in these ambulatory surgery centers. They were just doing them as non-owners. And so they're effectively taking physicians that are already operating in the facility and converting them into ownership, soon after private equity comes in.

Okay, so in the time that we have left, so now we're going to kind of shift gears to a wholesale private equity acquisition of a relatively large prominent ambulatory surgery center chain. Interestingly enough, as a bit of history, this chain called Surgery Partners started out as a firm in Florida in 2004, and is actually only in Florida up until private equity gets involved in 2009. So when HIG capital comes in, that's when Surgery Partner really goes for a national strategy. And currently, you'd see them in about 30 states across the US.

Also, important for our purposes is they're going to then do an IPO in September of 2015. And this is interesting in at least two ways, one of which is considerably different than what we're thinking about in the standalone individual firms, where they would not be engaged in an IPO. That when the private equity firms and the physician owners cash out or liquidate their positions, they're essentially selling them to other local healthcare companies, oftentimes, hospitals.

But here we're able to actually look at something that's a little bit closer to kind of the standard playbook for private equity we'd see in other industries and sectors of the economy, where the portfolio companies actually large enough to then go into public markets. And so we can see what happens when private equity takes over and what happens when they return to the public market.

So we get the traditional diff-in-diff setup. We're going to use roughly a decade of data to do this. So we're going to see several years before private equity comes in, the full duration of private equity, and then a few years after the IPO. And we have the simple setup, but we're going to take look at the surgery partner firms beforehand, and then again compare them to ASCs that never have any direct or indirect involvement in private equity, and were not part of the original treatment group in the prior analysis.

So similarly, just as what we saw in the standalone, the Surgery Partner's ASCs are much more tilted toward the privately insured market. They're also charging a lot more for their particular services, but that seems to be largely driven by case complexity. And then they also have quite a bit of physician owners, so about six on average. So that seems to be part of their financial model and their organizational structure at baseline.

But when private equity comes in, we don't see any obvious changes in case volume, so, again, not really a push toward productivity. And also really not a change in their kind of charging or list price behavior. Now things seem to tilt upward as they start approaching that IPO endpoint and beyond, which the data are not, I think, precise enough to draw too firm of conclusions, but at least it's suggested with the idea of trying to increase revenue and valuation at the time that they're going to the public market.

We, also, again, don't see any increase in the intensity of procedures per case or a change in the case complexity. So they're not changing how much they do or what they do, but they are losing some privately insured mix from their payer group. And so much like what we saw with the standalone ASCs that once private equity comes in, there seems to be some kind of loss of negotiation or contracting on the private side. Which, again, could be reflective of both on the ASC firm side of willing to lose some contracts, but also some private insurers not willing to negotiate with these firms going forward.

And then last but not least, we can look at the number of physician owners. And the data at least suggests that as the private equity firms are ultimately going to be cashing out with this IPO announcement, the physician owners seem to be kind of doing the same thing to leverage or cash-in their equity stakes.

So that was a lot of pictures and a whole lot of other pictures that I didn't show you for time. But to step back and do a brief summary, when we look at the clinical conduct, which again is one of the reasons there's a lot of attention and interest in private equity and healthcare, we're not seeing that the behavior of ASCs is radically changing. So it seems that physician decision making as well as their agency seems to be largely preserved.

What seems to actually change with the involvement of private equity is kind of the financing and ownership structure of these firms. So one of which is higher list prices, which again, just because of contracting structures can translate to higher transaction prices for services. But then, also interestingly, bringing in these physician equity owners and then coordinating their divestment decision.

One way to think about why they would want to do this, why they'd want to bring in additional physician owners, is it's a way to tie key human capital to these firms. And so a lot of the value built up in these firms is around the reputation and referrals that are tied to the physicians that do cases in these firms. And in many ways, you could think about those intangible assets as being worth quite a bit more than the tangible assets of these otherwise relatively small surgical facilities.

And this also kind of contrast with the private equity investments that we might see in other industries, particularly industries that have larger firms and a workforce that's maybe lower skilled. Whereas in these ambulatory surgery centers, almost every worker in there is high skilled, so it's nurses, surgical techs, physicians and so on. And so, again, being able to demonstrate to the entities that are ultimately going to buy these firms that the reputations there, the referrals are there, intact, adds a lot of probably credibility and negotiation leverage.

It's also just as we noted before that they're converting existing physicians into owners. Those that are then divesting with the private equity firms are not exiting these ASCs. They keep doing those cases here, which again, kind of, I think helps with this credibility story.

Now, that being said, when we looked over it on the chain side, now, it was just one chain. So in many ways it's just a case study of a particular horizontally integrated entity, but we don't see the exact same behavior. And so that at least suggests that there may be important heterogeneity based on the type of firm and the organizational structure, even before private equity comes involved. There may be some behavior related to the IPO, but again, the changes were not large, and it relates to this one particular context.

But the results are kind of consistent with private equity firms having these bespoke strategies for ASC investments. And so not just kind of walking in there and trying to figure things out, but actually having a good idea of how these outpatient surgery markets operate, what the contributions and value propositions of ASEs are, and doing what they can to, in some sense, improve those things.

And we also weren't finding evidence that was really consistent with things being harmful for patients as long as alternative providers are available. So the privately insured group, in particular, would be the ones that would be potentially most sensitive to higher list prices and higher transaction prices, because the government payers would be able to administratively set prices. But they were also the ones that were pulling back and not relying on these firms so much going forward.

And so long as these other providers are available and they don't have a kind of captive market, the data don't seem to suggest that private equity in the ASC industry space, specifically, is really leading to some loss of consumer welfare benefits. And I think that is all. Okay.

Speaker 2:

To discuss this paper, I believe, live from Boston should be Leemore Dafny. Leemore, are you there?

Leemore Dafny:

I am here. Can you hear me?

Speaker 2:

Yes, we can.

Leemore Dafny:

Okay, fantastic. Thank you very much. Thank you for having me. I really wish I could be there in person. I was teaching this morning, so it wasn't possible, but I'm delighted that people are gathering there, and really glad to have the opportunity to discuss this terrific paper. So I want to start by emphasizing, and emphasize it even more strongly than the authors themselves could, because it's their study that this is a really important topic. We're in the midst of a fairly significant national debate, or maybe I should say national handwringing about the role of private equity investments in healthcare services. And this paper contributes to the growing empirical evidence on the effects of PE.

So just a couple of basics for those in the audience who don't think about PE on a daily basis, as certainly some investors do. Private equity investors, they give a surge of capital to a business, in exchange for that they have an ownership stake and their goal is to turn it around, exit usually in less than 10 years, sometimes significantly less, and to earn high returns on that investment. Who doesn't want that?

There has been a surge in private equity investment in recent years. Just for one year of data, it was reported that there were 800 deals and a $100 billion invested, that's just in healthcare. So pretty substantial dollars going in. What PE firms do is that they acquire multiple practices, and where possible they issue debt, leverage up the company, and then sell their equity investment.

So there are concerns about this ownership form. Money is definitely being made. Money's definitely going in and coming out. But, of course, where is that coming from? So let me highlight a couple of the angles on this. First of all, clinicians are worried. I found this quote for an operator of an ambulatory surgery center, so I'll give you a moment to read it.

Okay. So this fairly well expresses you the concerns on the part of clinicians who are concerned that the new investors might interfere with their clinical judgment. Then payers, both public and private, as well as public, and by the public, I mean both policy makers and the average patient have raised concerns about private equity, too. So in particular, private equity firms have already consolidated a lot of formerly fragmented markets, ophthalmology, dermatology, emergency room physicians, anesthesiologist. And as you could tell by the specialties I mentioned, some of these are specialties that are renowned for having a surge of surprise billing.

And, ultimately, the No Surprises Act in the year 2020 has and will continue to address a lot of these issues with surprise billing. But not everything, notably ambulances have been left out, and there are so many loopholes in our healthcare system. And there is concern that private equity investors are likelier than the former owners of these practices to exploit these loopholes.

So there is a nice article written by a number of authors from USC, Erin Fuse Brown is the lead author, and I've given you the title here. I really like the title, Private Equity Investment as a Divining Rod for Market Failure, so you kind of know where they're coming from in that article. But they're not the only ones who are concerned, in particular, the chair of the FTC has expressed a lot of concerns, particularly in a recent enforcement or divestiture requirement and consent decree.

Okay, so this is too much for you to take in. I know that. This is mostly for you to look at later if you would like to, but in receiving this paper, I thought, gosh, in the last couple years, I know there have been a number of studies on the effects of PE acquisitions. I wonder how this fits in. And the authors helpfully cited a number of these studies. So here is just a quick summary chart, and I'm just going to highlight a couple of points.

First of all, I have two slides here. It does not list all the studies. I've focused just on event studies, so before/after analyses of acquisitions, and I have left off the studies on hospitals. And what you can see here is that the event studies on physician practices thus far have found that prices increase post acquisition. This is data from commercial claims with the transaction prices. And that quantities have increased. There is one prior study on ambulatory surgery centers, and that study focuses just on the Medicare population. They don't find changes in Medicare costs. Obviously, Medicare prices are fixed, so there isn't much to look at there with price. They also don't find change in volume.

So this study is going to focus on the entire payer population, and is also going to have a different study design than the preceding study. Just briefly, there are a number of studies on nursing homes. These studies use Medicare data and they find increased Medicare spending following PE investment, and also, largely speaking, reductions in quality homes using various measures, with some evidence that that effect is intermediated depending on the degree of market competition.

So I mostly just put this up here to say this literature is growing, and the more we know, the more informed we are. Also, we'll learn a little bit about whether the effects vary by the different sub-sector, and hopefully something about the drivers explaining why they might vary.

So what does this study do? That was a terrific presentation, so I won't have much to add on this. I'll just summarize for you. First of all, the study performs two empirical analyses and focuses on ambulatory surgery centers. There's only one prior study on ambulatory surgery centers, and it is limited in its lens. ASCs are fairly small operating facilities, as you heard, there are a lot of them nationwide, over 5,000. They are for profit. 90% half physician ownership stakes.

There are two analyses that we heard about. One is a differences-in-differences event study of the individual ambulatory surgery centers that are taking on their very first private equity investment with a very long post period in the data so the authors can see what happens. And they also follow what happens following divestment. They match to a sample of controls that don't have any PE investors, and in our counties that during that entire period don't have any PE investors either. And as they mentioned, they randomly assign a placebo date of PE investment in their controls.

Second analysis is a difference-in-difference analysis of what happens to large chain of about 21 ambulatory surgery centers in their sample, so in Florida. I know the chain is larger than Florida, but their sample has 21. What happens after that chain is acquired by PE investor, and then subsequently IPOed?

So the results are too numerous to go over here, but I'm going to highlight in the same colors as on the summary slides for the literature review, that the PE investments in the individual standalone ASCs are followed by increases in charges, which we believe are correlated with prices, or the increase in charges correlated with increase in price. And

Leemore Dafny:

... then a reduction in the share of patients at the ambulatory surgery center that have private insurance. So that is consistent with what we know from other studies of private equity-backed practices, certainly in the emergency room physician literature, that the private equity-backed practice decides to go out of network and then negotiates a higher contracted rate.

And that could lead to a reduction in the privately insured patient share and, ultimately, an increase in the price per case that they see. Okay.

On the second analysis, the author, after a chain, transfers ownership to a PE investor. There isn't any immediate effect on price. There is an immediate decline in private insurance shares, suggesting that the commercial insurers, after that PE acquisition, either steer their patients away actively or take the provider out of network. And post-IPO, that price goes up; the volume goes up.

All right. So let's talk a little bit about... I'm not sure how much time I have. That's an advantage of not being there in person because nobody can ping me. But I will try to speak quickly.

So just a couple of empirical comments, and then I'll talk about how to interpret the analyses. So one question is just trying to understand why the focus is just on the PE investment into standalone ASCs. I can see that that's how you might start. But it's certainly possible, I would think, that there are multi-site ASCs that are then rolled up over time. I just wanted to understand how common that scenario is and why it'd be left out.

The second point is one that I hinted at when I described the empirical strategy. The treatment group, those that receive a PE investment or take on a PE investment, are pretty different in observables from the control group. They are lower volume. The treatment group has a much higher commercial insurance share, among other things. They seem to be in different markets, certainly.

So one question, of course, is, is it possible to find controls that look more like the treatment? Do some propensity score matching. Barring that, or even with that, if we could see not just the difference in treatment relative to control, but the absolute trends in treatment and control, that would be interesting to try to understand what is happening in the markets where these different ASCs are operating.

Second point is taking a look at heterogeneity of effects. Okay. Of course, it's a small sample, and it will be difficult to get anything statistically significant. But it would be really interesting, even directionally, to try to understand if the effects are more or less pronounced depending on investor characteristics; in particular, whether that investor has a presence already in the local market or has some market power in other local markets.

So is there potentially some cross-market power being realized and exercised? Is there horizontal consolidation and price increases associated with that?

Relatedly, does what the private equity owner do after the acquisition affect the outcome? So if they roll out pretty rapidly, is that different from if they are acquiring in a fairly slow pace if you have enough variation in that?

On the second analysis, same comments on the control group. Here, again, the treatments are very different from the controls certainly in observables. And it would also be interesting to look at heterogeneity of effects, in particular, by specialty.

There's something pretty interesting in the results where the prices increase or the charges surge, especially for all other insurance, which made me wonder if there was something going on with orthopedics and workers' comp. Or is there some loophole, basically, that is being exploited there that we're learning about from that analysis?

And a suggestion to take a look at what happened to the debt following the PE takeover because that is a pretty common strategy to issue a lot of debt. And that leverage increases the payoff to the investors. And as we've seen, we've seen that the physician owners realize that payoff upon exit.

Ah, this won't be lost on the authors. Where are the regression results? How many times do you repeat the randomized date for the placebos to have hypothetically received PE investment? What are the standard errors? Actually, I see the graphs. I didn't see it in the appendix. Maybe it's still in the works.

All right. The fun stuff is how do we interpret all of this? And the authors say that it seems like there's not much of a change in clinical practice. It's financial engineering. And to some degree, that kind of comes as a relief.

But I want to throw out this question of, well, what enables that financial engineering? In particular, I'm not talking about favorable terms for borrowing. But what enables the price increase? It's not possible or it wasn't done. Anyway, it would be really hard to do, so I sympathize.

To look at what happens to quality, is there some reason that these services have become newly superior that it is worth paying more for them? Or is it the case that there are different negotiating tactics where the practices are more willing to try going or threaten going out of network in order to obtain higher prices, or they have more market power and are leveraging that. I do want to note that all the cost efficiencies that are possible should theoretically incentivize lower prices, which we don't see.

Second question is, okay, in the short term, no clinical changes, no changes in clinical activity, but what are the long-term effects on consumers of this transition? One question is, what happens to the labor market for these clinicians who are operating?

After the initial wave of physician owners cash out, then the NPV of wages and equity for the owners on an ongoing basis, who now are employees and a highly leveraged employer, does that reduce their earning prospects and lead them to want to be likelier to go get employed elsewhere and, therefore, cause labor costs to go up and possibly, ultimately, the exit of these practices?

On prices, you see that prices go up and that privately insured patients are less likely to utilize the facilities in due time. But how long can that last? Ultimately, will the private insurers end up caving and having their patients go there, or will they potentially end up with access issues? There's a question about the sustainability of that. And relatedly, the long-term competitive effects.

To the extent that this financial strategery ends up enabling investors to capture value and, hopefully, to create some of it, there is incentive to continue growing through acquisition, raising the prospects of both same and cross-market power. And thus far, no evidence of benefits passed on, at least in the form of price.

That's where I'll end. I apologize if I've gone over. And I just really want to thank the authors for giving me the chance to read such an interesting study and think about an important subject.

Speaker 3:

We have time for maybe one or two questions from the audience.

Ginger Jin:

Hi. Ginger Jin from the University of Maryland. Thanks for a terrific paper. I want you to give some comments on the relationship between price and quantity.

If I understand your paper correctly, it seems like the price is going up, but the quantity does not change. Does that mean somehow patients are not price-sensitive or there's induced demand or they're stealing patients from nearby ASCs?

Michael R. Richards:

Yeah. Normally, I speculate wildly, but this is being recorded, so I'm going to be careful. One way to interpret that is, again, thinking about the negotiation between the insurer and these facilities, but also recognizing the fact that among the mix of payers that they have, some of them are largely insensitive or unexposed to those list prices because they set their own prices.

So, Medicare, traditional Medicare being the obvious example, it tells ASCs, "This is what you get paid for what you do." So it's really the privately insured group that is most exposed. There are obviously other smaller payers that would also be engaged in negotiation.

One way that we're thinking of it, but I think Leemore has brought up a terrific set of ideas of other ways to think a little bit more about it and enrich the understanding, but this idea that they effectively say, "This is what we would like to be paid." And then they're going to go back and forth with the insurer.

And if the insurer says, "Well, we're going to traditionally contract 60% of prices with you," 60% of a big number is better than 60% of a small number. But we do see this drop-off in their payer mix that's pretty substantial and happens pretty quick among the privately insured group, which is suggesting that either those contracts and those negotiations are breaking down, or as Leemore pointed out, they're changing the way that they're networking and steering their enrollees toward those facilities.

David Benson:

Hi. David Benson, Federal Reserve Board. I was wondering if you could say anything given your data, not knowing your data limitations, about more economic-oriented measures of the physician's practice.

So, clearly, measures of efficiency, like how many services for a patient, things like that don't seem to be changing. But what about the capital-to-labor ratio or the intensity of materials going into the types of procedures? And then maybe that would shed light on how these practices are using the capital that they're getting.

Michael R. Richards:

Yeah, I agree. I think that's a great point. Especially when you think about the standalone ASCs, these really are small mom-and-pop firms.

And so the arrival of private equity could have some management implications, in particular, when it comes to negotiation with contracts and so on. But another element to it could be doing a lot of these back-office functions, so getting better prices on inventory and supplies, doing a better job with health IT and revenue cycle management, and these sorts of things.

I think all of that is possible, and we absolutely want to look at all of it, except the data do not exist; whereas in the nursing home industries and the hospital industries, we often have at least some data in those domains. When it comes to ambulatory surgery centers, the data are really, really hard to come by.

I mean, even to this day, MedPAC has been engaged in a back-and-forth fight with them about quality reporting from ASCs, which is something that has historically not been a part of the ASC industry. So I completely agree.

It's just we have not found... even going through other entities outside of the traditional data spheres, to get that type of information. But we are constantly searching. And if anybody comes across, we would love to see it.

Speaker 3:

Great. Thank you very much for everyone's participation in that session. Now we'll introduce the second paper.

Michael R. Richards:

And thank you, Leemore, if you're out there.

Speaker 4:

Hi, everyone. Next, we're excited to have Zhenling Jiang from Wharton School presenting her paper on designing dealer compensation in the auto loan market.

Zhenling Jiang:

Thanks, Will. I'd like to thank the conference organizer for giving this opportunity for us to share this paper with you guys. So now we'll transition from going to hospitals or ASCs to get surgery to going to auto dealers to buy cars, probably equally pleasant for consumers.

Now I know why everybody is getting confused about what to push. Let's see. So first, I'd like to recognize my amazing collaborators in this project. Max Wei is from USC. Tat Chan is from WashU. And Naser Hamdi is a VP of analytics at Equifax.

We know that in many markets, the transactions do not happen directly from the consumers to manufacturers, but are facilitated by this intermediary or this middle person. And in such a process how this intermediary or middle person gets compensated can influence the transactions and influence the final consumer outcome.

So we are looking at this market of indirect auto lending. And in this market, the auto dealers, they are this middle person. They're connecting consumers who want to take loans to finance the car purchase to the auto lenders who are supplying the loans. And in particular, we're interested in how dealer compensation in this market will influence the final consumer outcome.

So why do we care about auto lending market? First of all, I just checked last night, the most recent stats says that US consumers hold $1.4 trillion in balance in auto loans. Obviously, a very, very large and significant market. But also, it's also a significant portion of profit for dealers.

So if we think about auto dealers, then we may think that, okay, they're in the business of selling cars. But if you look at how dealers are making money, the story is more so that they are in the business of selling cars so that they can make money from financing and insurance that goes with a car purchase.

Increasingly these days, even a much larger proportion of the total profit actually comes from the profits coming from the financing part. So, obviously, [inaudible 03:12:08] important issue for auto dealers to care about as well.

So just as not everybody studies healthcare, I also realize not everybody studies auto financing. So how does it work? As a consumer, we go to auto dealers. We buy cars. Many consumers, over 80-85%, also get financing at the dealership.

So how does that work? After we talk to the salesperson to nail down the cars and the car price, we talk to this finance manager at the dealership who have connections with maybe multiple lenders. And they'll get a quote from the banks depending on, say, our creditworthiness.

And eventually, the loan contract is from the consumers and the lenders. And the consumers will make a payment to the lenders afterward.

So how do dealers make money in such a process? The auto lenders will give a quote depending on the consumer's creditworthiness and what type of loans the consumer is asking for. So for example, the bank says, okay, given this consumer, I'm asking for 3%. And we call this the bank receiving rate. So this is how much money banks eventually make from this loan transaction.

What a lot of people don't realize is that after that 3%, dealers will stack another dealer markup (in this case, 1%) on top of that price. So, eventually, what consumers end up paying depends on both this bank receiving rate as well as this dealer markup.

Also, super important to know in this context is that in most cases, the dealer markup is discretionary, meaning that each consumer, each transaction, gets a different markup. And this markup goes towards the dealer compensation. So you can imagine dealers have incentive to really increase this markup for consumers where they think they can make the money.

So, obviously, if we think about this market, there's really room for discrimination in pricing. Different consumers based on the same credit scores will end up paying a different price.

So one of the most important law governing the consumer lending space is this Equal Credit Opportunity Act. It says that lenders cannot discriminate based on things like race or gender or age among other things.

But if we look at the structure of this market, because this markup really varies by transaction and is not based on the consumer's creditworthiness, then there's room for a certain type of consumers, in particular, consumers who are more disadvantaged, say, like minority consumers, that they will pay a systematically higher price than other consumers based on the same credit profile who are applying for the same type of loan.

So this is exactly the issue that regulatory agencies have been looking at. In particular, CFPB and the DOJ have fined several lenders for this alleged discrimination in pricing. So they're saying that, based on the same creditworthiness, we are seeing that minority consumers end up paying a higher price than other consumers.

So I'm giving one example here, but there have been a series of lawsuits around this time that they're going after lenders for this practice. Now you may be thinking, okay, lenders are not really discriminating. It's the dealers who are adding this markup. But ultimately, what CFPB says is that it's because as a lender, you're allowing a dealer to add this markup, and that's the dealer compensation.

So what the CFPB is advising lenders to do is to say that, okay, you can still compensate dealers, but instead of having this discretion to markup, we are advising people to fairly compensate dealers using another mechanism. In particular, have a flat fee per transaction so that there's no room for discrimination.

So as you can imagine, those years in this type of regulatory environment, there are several lenders who voluntarily change the practice in compensating dealers. In particular, they switch from this discretionary markup to a fixed way to determine the markup.

In particular, they're saying that we're just going to pay a dealer 3% of the loan amount. So then the money is going to be deterministic. Actually, based on how large the loan is, I know how much I'm going to compensate the dealer.

So with that fairly long setup, what do we do in the paper? The paper has two parts. And the first part is going to be descriptive. We are just going to document after these lenders make this change in the compensation of dealers. How does it impact consumers? What's the impact on lenders?

The second portion of the paper is going to be prescriptive. We're saying that, knowing how this market works, we're going to want to design a different type of dealer compensation. And we're going to take the bank's perspective. We still want to ensure this non-discretionary compensation, but we want to improve the bank's market share.

As a preview of the results, what do we find? We find that the consumer rate indeed decreased for these more disadvantaged consumers. So this is in line with the intended policy goal for these banks. However, we show that, interestingly, the market share for these consumers, for these banks that changed the policy, also decreased.

We're going to use this as our key reduced-form evidence to show how dealers also have a significant involvement in such a market in determining what type of loans consumers will eventually get. We're going to build on that reduced-form evidence to build a bargaining framework to characterize the incentive from both the consumers and the dealer in this market.

And based on that, we're going to be able to run our counterfactual policy to say what if the lenders can choose a different type of compensation, and how is that going to work? What we'll show is that if lenders go with a fixed lump-sum policy, that's going to both ensure consumer protection, but at the same time, improves the lender's market share.

Okay. So what data do we use to study this question? We get loan-level information from these two groups of banks. The first type, the target bank, these are the banks that actually made this policy change. So everybody used to have discretionary; the dealer can do their own thing. And now these target banks switched to say we're going to pay you 3% of the loan amount.

We're also going to get loans from what we call general banks. But these are just control banks who maintain their previous way of doing business. We're going to get loan-level information. And for each loan, we're going to be able to see the loan-level characteristics as well as some consumer characteristics. And in the end, there are about 180,000 loans in our sample.

So first off, let's see what happens on consumer rate after they have this policy change. So over on the left side, I have this... This is just coming from raw data and running a density plot. It says that the distribution of how the consumer rate changed a lot after these banks changed their compensation policy. And this is quite intuitive.

If we think before the policy change, then because of this discretionary markup, then there's a lot larger variation of the price coming from some consumers getting a really good deal at a dealership, so they're able to get a low price. But some consumers are really getting screwed and having a very high price.

So after they change from discretionary to this fixed way of compensating dealers, then that sort of variation goes away. So then the price for a consumer gets a lot more concentrated after the policy change. So this really makes sense, looking at the target banks. For our general banks who maintain their usual way of doing business, during the same time period, nothing changed.

Now I'll talk about perhaps one of the most interesting reduced-form evidence we show in a paper. What is this doing? So we're dividing consumers into different groups. So as I was saying, previously, some consumers, if they are, say, disadvantaged, they're really getting screwed over by dealers. These consumers should benefit.

So now they have this fixed way to calculate markup. And indeed, now on the left side, I'm breaking consumers based on their credit scores. And what the graph is showing that if you are lower credit score consumers who tend to be less sophisticated in terms of financial decisions, we are seeing a decrease in their price. It makes sense.

What's interesting is what happens to the market share for these banks after the policy change. As I trust that we all know from ECON 101, the demand curve should be downward sloping. When you have a lower price, I should see a higher demand. And this is the opposite of that downward-sloping demand curve.

So now I'm showing that when a group sees a lower price, they also have a lower demand. So what's happening now? If this was a seminar, I'm going to invite people to guess. But since we are in a larger audience, I'm just going to give the answer, which is the auto dealers.

So consumers is not in like a supermarket setting; I see all the options and I'm going to pick one. In this market, dealers are working with consumers to present this loan option to a consumer. So what happens is that when, say, the dealer is seeing a consumer they think is really naive, I can charge them a very high price.

I'm not going to give the loan for these target banks, even though it's going to benefit the consumers with a lower price. I'm going to steer those consumers to the other banks where I can still make this discretion of markups. We still have these control banks who allow for having as high of a markup as they want.

So even though consumers would benefit in such a case, but then because of the dealer involvement, these banks are actually seeing a lower market share. So a similar logic goes for... you can divide consumers in different ways. Right? So now instead of saying credit scores, now I divide consumers based on where they live. It depends on a minority percentage. So very similar stories.

I see that, for consumers living in high-minority regions, they are seeing the largest decrease in price. We are also seeing the largest decrease in demand. So these consumers would benefit from this new policy, but they are steered away by the dealers to these other banks. Okay. So I hope that was clear.

Building on that, we're going to take this learning that we see from our reduced form, and we're going to characterize the market with a structural model. So what's the key component going into the model?

We are going to take consumers and their demand for loans as given. So I'm a consumer, I want to borrow 20,000 to buy a car. The banks, in our setting, they're going to give this bank receiving rate and this quote depending on the consumer characteristics; in particular, say, credit scores, the type of loans that consumers are applying for.

And then they give this information to the finance manager at the dealership. And then eventually, we're going to characterize the interaction between the consumer and the finance manager with this bargaining framework. In particular, we're going to use Nash bargaining.

And we're going to say that your final output, so this final deal that you reach, is going to depend on both the incentives of consumers as well as the incentives coming from the finance manager, coming from the dealers. And then people are going to vary based on the level of bargaining power you have in this transaction.

So we're going to leverage that. And we're going to use the model to characterize the eventual price, consumer rate, that people get as well as which bank is going to end up financing the loan. Is it the target bank or the general banks?

I'm not going to go into too much detail on the model, in particular, before the policy change. We have the general banks and target banks. They both use this discretionary markup.

So things are fairly straightforward now. It's just that the consumer price, the final consumer rate, is going to be varying depending on the bargaining power for consumers. So given the same type of consumer characteristics, same type of loan demand, some consumers get a higher price, some consumers get a lower price. And it's going to be determined by the bargaining power of a consumer.

What's more interesting is what happens after a policy change. So now, dealers are faced with two types of banks with different compensation policy. So they can either go with the target banks that have deterministic ways of compensation, that is 3% of a loan amount, or they can go to these general banks who still maintain these previous ways of compensation.

So now if you think about it, which banks that they prefer can be different. So the banks that benefit consumers or the banks that a dealer prefers now can be different.

So now in this after-policy framework, which bank ends up financing the loans is going to be also... it is a bargained outcome. So it determines... It's going to be dependent on the incentive of both parties as well as their bargaining power.

Okay. So one of the most important parameters that we're going to estimate is this bargaining power parameters in the model. So without going into too much details, I'll highlight the identification of this bargaining power parameter.

In particular, in our model, in our data, we do not observe the bank receiving rate. So we don't really know how much the bank is quoting the consumer. So eventually, we're going to rely on this policy change and the relative change in market share to pin down this bargaining power parameter.

So what do I mean by that? This is very similar logic as what I showed you in the reduced form. So for example, if I see a group of consumers, so for example, this low-credit-score consumer, I see that they would benefit a lot from going with the offers from the target bank after a policy change. However, their market share after a policy change coming from the target bank also decreases a lot.

So knowing that, I'm going to infer that these consumers must have low bargaining power because they would benefit a lot, but then, in the end, they're not getting these better loans. So they're now inferred they must have lower bargaining power.

So this is very similar logic as what I showed you in the reduced-form evidence. And this is going to be the key thing that we're going to utilize to pin down what type of consumers have low or high bargaining power.

So what do we see? We estimate the whole model and highlight our

Zhenling Jiang:

For bargaining power parameter estimates, we see that consumers' bargaining power tend to be higher for consumers who have a larger loan amount, and we think this is quite intuitive. The loan amount is most likely to be determined based on the car price, so consumers who buy more expensive cars, getting larger loans also tends to be kind of like more resourced and more financially savvy.

We also see that bargaining power tend to be lower for consumers whose loan term is longer. Also makes sense. They are probably more constrained on the monthly payment, so they have to stretch the loan out longer. We also see that consumer who have a higher credit score tend to have a higher bargaining power. Also quite intuitive. We also estimate the variance of this error term in bargaining power, and that's going to say, okay, beyond all of this, there's also some variation that is unobserved in the data that contributes to the difference in bargaining power across the population.

After getting the results, we can do one thing interesting, is for each long observation, we can back out this unobserved portion of the bargaining power in the data. So, based on your eventual loan outcome, based on your rate and which bank you're going with, I can get an estimate of your [inaudible 03:31:36] loan term in the sample. So, after getting that, we can do something interesting, right? So, now I have this estimate of the bargaining power residuals coming across population. I can relate that to the zip code level demographics. So, here, we see that if you are consumers coming from a higher minority region, surprise, surprise, they tend to have lower bargaining power. So, these are the consumers who, all else equal, will get higher price, who are more likely to get loans that are more like less favorable to them.

So, after showing that, we are coming to our second goal of the paper, which is to design compensation policy on behalf of lenders. We want to maintain this non-discretionary in our framework in the sense that we no longer allow dealers to have this room to determine how much markup I'm going to impose for each consumer, and instead, we're going to have a deterministic way to find what is the markup.

So, in particular, we consider three type of policies. The first one is going to be a fixed percent of loan amount. So, this is consistent with what the lenders are deployed in the market. The second type we consider is how to have a fixed dealer rate. This is sort of close to how they're doing things now. They're adding a percentage point on the loan, say like 1% or 2%, but instead of have that percentage vary across individual, now, we say just the fixed number that applies for all the loans. The last type is a fixed lump sum. So, we're just saying that simply for each loan, you get paid this much as your compensation. This also happens to be the recommendation of CFPB back in the days when they published that bulletin, and I'm happy to report that that is also the best performing that we find in our counterfactual.

So, when we compare across these three types of compensation, what we find is that if lenders are compensating dealers by just giving them a fixed lump sum, the lenders are able to achieve a higher market share compared to their adopted policy, which is compensating dealers by this percentage of loan amount. So, now, if we think about why is that, that seems a bit counterintuitive, right? So, for lenders, you may think that if I'm rewarding dealers by giving them a percent of loan amount, I may incentivize the dealers to get me more like larger loans. I like to have larger loans. Maybe I'll just compensate based on a percent of loan amount. That seem to be intuitive, but now, if you consider the bargaining power and how also both the consumers and dealer plays a role in this market, that is actually counterproductive. Why is that?

So, you want to reward consumers in the case when they have a high bargaining power. I just showed you that people who have larger loans tend to have higher bargaining power. So, those are exactly the cases where the lenders want to give consumers better price, and you want to pay dealers roughly less. So, compensation by percent of loan amount is actually exactly counterproductive from that perspective.

So, why does fixed lump sum work so well? The reason is if we fixed the amount of compensation, then in terms of this percentage, that percentage is going to be lower if my loan size is bigger, right? So, I'm just giving you 500 bucks that relatively, percentage speaking, is a lower percent if I'm talking about a much larger loan. These are exactly the consumers we know have high bargaining power, and those are the consumers you actually want to give dealer less and also offer consumer a better deal.

Okay. So, what do we find in this paper? In this paper, we study this question in the indirect auto lending market where the lenders need to go through dealers to reach the final consumer. We document this, what we think is very interesting case of invert, like the reverse of demand curve. When price is dropping, demand is also dropping. Why is that? Because it's not posted price market only chosen by consumers, there's also this middle person, this intermediary, the dealers. Also, their incentive matters in this loan process.

So, to model things up, we apply this Nash bargaining framework to model this joint decision making between the consumers and the dealers, and we're able to find that having this fixed lump sum compensation, consistent with the intuition of CFPB, is actually works better than the adopted policy from the lenders. We believe that this paper have some implications in this important market of indirect auto lending. In particular, it's important for both the lenders and regulators to consider that in this market, whenever you are designing policy, it's extremely important to take the consideration of the incentive of dealers because all the things you're designing in terms of the consumer-facing price, they all need to pass through dealers to reach the final consumer. If the things you're offering is going to be going against the incentives of dealers, it's not going to be able to effectively reach the consumers.

So, in particular, we show that it totally makes sense for the dealers, for the lenders to compensate for larger loans but actually goes against what we find in terms of these consumers have higher bargaining power. Therefore, you actually want to do the exact opposite of that. So, and from a methodology perspective, we extend using the Nash bargaining framework in terms of a demand estimation issue. We believe that this framework can also extend to other situations where the firm price needs to pass through this intermediary to reach the final consumers, and this intermediary have their own incentives in place, and it's important to consider both parties when designing the estimation system. Okay, I'll stop right here. I look forward to hearing our discussion and also your questions. Thank you.

Speaker 5:

Great, thanks. Now, we have Pranav Jindal from UNC Kenan-Flagler Business School to discuss, and he's going to be joining remote.

Pranav Jindal:

Perfect. Thank you. Can everybody hear me? Yes? Okay, great. So, thank you. Thank you for this opportunity to discuss this fantastic paper on designing legal compensation in the auto loan market. So, I'll jump right in. Zhenling, thank you for a great presentation.

So, I'll give a brief overview over here, what I would say is a helicopter view. The paper's about designing dealer compensation in the auto loan market. Just the loans itself, $150 billion market growing at about 8% per year, so it is an important market, an important and timely topic to study. Also, barring the last three, four years, there's been prosody of empirical research in on bargaining in the B2C domain where retailers, sellers are negotiating with end consumers, so the paper fits in very nicely into, what I would say is still an upcoming stream of empirical research, and the findings are very managerially relevant and well-motivated.

As Zhenling mentioned, a little bit of an overview over here, traditionally, dealers would add a markup on top of the bank's recommended interest rate or the bank's receiving interest rate. This resulted in discrimination where people with different credit scores, but not only credit scores, based on different demographics were paying different interest rates. So, policymakers advocated for non-discretionary compensation schemes. In these non-discretionary compensation schemes, we can hold fixed either the dealer's compensation either as a percentage of the loan or the interest rate, or there could be a fixed lump sum payment.

The key variation that Zhenling and her co-authors employ is a switch in policy for the target banks where they go from discretionary to non-discretionary scheme with a fixed 3% commission. The key figure in the paper, and Zhenling discussed this as well, is the relationship between interest rates and market shares for the target banks before and after the policy change. So, I'll spend just a minute over here. On the left, we see the interest rates that consumers pay at the target banks before the policy and after the policy. What we find is that the interest rates are lower for consumers who have lower credit score. As we go from left to right, the credit score increases, and the interest rate after the policy change increases.

So, looking at this picture, because interest rates are going down for low credit score consumers, we would expect the market shares to increase, but that is exactly the opposite of what we observe. That is the right figure where the market share also goes down among the low credit score individuals after the policy change. This reversal of demand curve is what, in the paper, is attributed to dealers and the role that the dealers play as intermediaries between the banks and the consumers.

So, in terms of the nuts and bolts then in the paper, the interest rate is determined based on a negotiation between the dealer and the consumer, which is modeled as an outcome of Nash bargaining. The authors are very careful to recognize that not all the variation in the interest rate could be coming from negotiations, but a substantial chunk of it can be attributed to individual level differences. In addition to the interest rate, the choice of the bank is also negotiated between the dealer and the consumer. So, the way to think about it is that once the interest rate is negotiated, we can determine the surplus that the dealership gets. We can get the surplus that the consumer gets, and now, we compute some sort of a weighted sum of these two utilities or these two surpluses where the weights are again the bargaining power.

So, the idea over here, as Zhenling alluded to, is that if a dealership has a higher bargaining power, then not only will they be able to bargain a higher interest rate, but also they will be able to steer consumer to a bank, which is more preferred by the dealership. The model then is estimated using methods of moments, and in terms of counterfactuals, there are three policies which are considered, and each of them, we are holding fixed something. It is either the percentage of the loan amount, the dealer rate, or the lump sum payment. As Zhenling mentioned, the lump sum payment scheme results in highest market share and consumer welfare. The underlying intuition here is that it best aligns the dealer's rate with the consumer's bargaining power. So, again, very intuitive findings. Great paper. Great analysis.

In terms of my comments, I have some thoughts, more clarifications but wanted to put them out. I will break them down along three dimensions. So, a little bit about the institutional details, the model and estimation, and some suggestions on the counterfactual analysis. So, starting with the institutional details first, when I was reading the paper, one question, which came into my mind was whether the dealership has to disclose all the interest rates to consumers or interest rates of all the banks. What about a situation where the dealership gets the interest rates from the banks? Let's say both target banks and general banks has some expectation of what the negotiated interest rate would be and based on that, only discloses one interest rate to the consumer. So, the consumer then is basically just choosing whether to take that interest rate or not. The dealer's role then is not just about negotiating but also about selectively disclosing information. This has implications for how we think about the model, but some light over here would be very helpful.

Also, in the paper, there is an assumption that the loan term and the loan amount are predetermined, and the interest rate is negotiated conditional on this. I understand the attractiveness of this assumption, but it would help to see some evidence to rule out the possibility that these could be jointly determined. So, for example, as a consumer, I might be willing to take a higher loan, make a smaller down payment if I can get a more attractive interest rate. So, these could be simultaneously determined. So, some regression where, let's say the interest rate is regressed on the amount of loan term or the other way around, controlling for observables could be done or could be run to test this assumption.

Going to the model then, there are three pieces I would like to briefly discuss. The first one is the consumer's reservation rate or the interest ceiling, which a consumer would pay in the model that is RI. Typically, when we think about the Nash bargaining, the surplus is defined based on the consumer's reservation price, which would typically, in this case, be their willingness to pay and the dealer's reservation price, which would be the cost from the bank or the interest rate from the bank.

In this case, the willingness to pay is replaced by this reservation rate or interest ceiling RI. In that sense, it acts more like a posted price, which also implicitly implies that the consumer's willingness to pay is higher than this interest rate. Now, that in of itself is reasonable to assume, but then that same reservation rate RI is also used in the counterfactual. So, it's treated as a structural parameter, which is invariant to this change in policy, which is where some clarification would be helpful.

The second thing I would like to get more clarity on is the interpretation of bargaining power. As I mentioned earlier, bargaining power typically represents a lot of other primitives, which could be the cost of going back and forth, impatience, and so on and so forth. The Nash bargaining solution is a reduced form approach to capture all of them and split the surplus. In the paper, however, not only is the interest rate and outcome of the Nash bargaining power but also the bank choice is a function of the bargaining power. So, how should we really think about this bargaining power and also the identification of bargaining power that is coming from both the observed interest rate as well as the observed choices?

Finally, in terms of the modeling ... and this is a point, which I don't have a very good handle on. I've debated this myself as well. What we observe are only the negotiated prices of the chosen alternatives or the chosen banks. So, we don't observe the negotiated prices for the non-chosen banks. Now, all else equal, we can reasonably assume that the non-chosen banks had worse negotiated interest rates, so they were more unattractive. So, to put it differently, we probably had some bad draws for those prices. So, how should we take that into account? How should we model them? I think in the paper, the model is based on the expected negotiated rate. So, it would help to understand a little bit more about, is this expected or maybe a outcome, which is worse than expected.

The last part that I want to say something about then is the counterfactual analysis. So, the analysis right now assumes that there is no response from general banks, and I think that is fairly well-motivated because the authors first show that during the duration of the data, there is no change in the interest rates from general banks. There is no response from them, so that is great, but having said that, if you think about the medium to long run, any change in policy by the target banks could also generate some sort of responses from the general banks. So, it might help to think about competitive reaction in this setting especially if the policy is not uniformly implemented across all types of banks.

Finally, just building on the counterfactuals, and this is maybe just wishful thinking, but could a bank specify the non-discretionary compensation differently for different consumers? The difference over here could be based on credit score. For example, the commission rates to the dealership might vary, or the amount of lump sum payments might vary depending on the credit score, or to put a completely different hat on, the payment mechanism itself, whether it is commission rates or lump sum payments or some percentages, they vary depending on the credit score.

So, just some thoughts over here to augment the current analysis. To summarize then, this paper presents well-executed analysis of an important and understudied area. It's a big market, very important. The analysis leverages variation in policy, which is great because even in the absence of model, we can see variation in the data, which speaks to the role the dealerships play. Finally, the findings over here are very well explained. They're managerially relevant, so I congratulate Zhenling and all the authors on a great paper, and I believe the paper just got accepted at Marketing Signs. So, congratulations on that as well. Thank you once again.

Speaker 5:

Great. Thanks everyone. We have time for just a couple questions before break so ...

Ginger Jin:

Yeah. Ginger Jin from University of Maryland. A fascinating paper and congratulations on the publication. I'm sort of intrigued by two institutional details here. One seems like consumers don't have full information about the bank receiving rate and how much the dealer is getting out of this, so I wonder whether some disclosure regime there forcing the dealer to disclose those two separately would be helpful. Not sure whether your paper can say anything about that and-

Zhenling Jiang:

Yeah. So, what I can say is there is no disclosure requirement for consumers. I trust that many of us have bought a car at a dealer before. I challenge you, none of you know what this bank receiving rate is. This is a very opaque market, and Pranav is exactly right. So, dealers don't really show consumers, "Oh, I get this three quotes, and this is what they look like. This is how much I'm charging you."

Dealers will be crazy to do that. So, I've talked to several dealers, finance managers in particular when working on this project. What they're describing to me as the process of them setting the rate is, okay, they'll say consumers walk in. Then, you talk for a long time at a dealership. They will get a sense of who these people are. If you are a professor in Maryland, you're probably very smart, I'll probably give you a better price. If you look kind of more naive, like more trickable, I'll give you some other price. So, this is as if they are bargaining, and the outcome will vary based on these individual characteristics, but consumers don't really know what the banks is charging them.

Ginger Jin:

Okay. Yeah. Thank you. My second question is, seems like the CFPB policy has created a non-level play field between the target bank and the general bank and, in this case, seems like to the detriment of consumers. So, I wonder whether you have any thoughts on sort of, should this policy roll out to all banks, requiring them all have the same non-discretionary and requirements or even just maybe a simple disclosure whether this rate was coming from a target bank under this non-discretionary versus a general bank with the discretionary would help?

Zhenling Jiang:

Yeah, that's a great thought. First, I need to clarify. There is no regulation coming from CFPB that you need to move from discretionary to non-discretionary. So, they are advising lenders to do so. I believe back in the days, if you read the news articles back around those years, then you see a lot of discussion both on the lenders' part as well as the dealers' part. They are anticipating something like this may happen, but as we all know, that the auto lending market is very political stuff. So, after a change of administration, then this policy also kind of shifts with that. So, even at the moment, there is no regulation coming from CFPB in this market. I was told that there are some folks from CFPB today in the audience, and I would love to talk to you more privately afterwards, but great question. Thank you.

Speaker 5:

Great. Thanks everyone. I think that's all the time we have, and we'll reconvene here at 3:00 after another break. Thanks.

Speaker 6:

Hello. Hello everyone. We're going to get started. We're not? Good afternoon. We're delighted to have Julie Holland Mortimer who'll be giving us a keynote address on diversion in the use of second choice data. Julie is the Kenneth G. Elzinga Professor in Economics and the Law at the University of Virginia. Prior to joining the University of Virginia, she previously worked at Boston College and Harvard University. Julie.

Julie Holland Mortimer:

All right. Thanks very much for having me. I'm unreasonably excited to be at an in-person conference. So, what I'm going to talk about today ... Let me see. Let's see. Oh yeah, there are my slides. So, everything that I'm going to talk about today is joint work with Chris Conlon at NYU Stern. I'm going to talk, sort of going to give you an overview of some previous results and then talk about some new work that we're doing with a student at Boston College, Paul Sarkis who is our co-author on that work.

So, let me just start off talking about diversion ratios. So, diversion ratios are what I think of as the best way that we have to quantify the strength of competition between products, okay? So, here's the idea. You raise the price of good J and then you count the number of consumers who leave, right, and the diversion ratio is the fraction of those leavers who switch to K, and that's diversion from J to K. So, it's this pairwise concept. So, a higher value of diversion, of that diversion ratio indicates sort of closer measure of substitutability between those products, so more competitive in that sense.

So, one of the things that makes this really useful is that it arises naturally out of sort of a multi-product, Nash Bertrand differentiated products model. You can see on the slides, there's the diversion ratio is this opportunity cost. So, marginal revenue equals to marginal cost here. This is sort of the opportunity cost of increasing your price of product J is that people are going to leave, okay? All right. So, it's helpful in merger context because if we internalize that diversion, that's going to lead to an incentive to increase price, okay? So, we can also write down diversion as the demand, as the ratio of demand derivative, so the derivative of the demand for product K over the derivative of the demand for product J. That also allows us to think about it as the ratio of cross to own-price elasticities multiplied by those quantities.

So, there are a couple of big advantages to diversion, and diversion was initially featured ... I think it goes all the way back to the 1992 Merger Guidelines but then really was in the spotlight and in the 2010 Merger Guidelines that Joe Farrell and Carl Shapiro took the lead on, and sort of helps us really to connect to differentiated products markets much more closely and also ... Sorry, and also has a number of other advantages that I'll talk about here for a minute.

So, one of the nice things about diversion is it allows us, as opposed to cross price elasticities, they're all normalized, right? They sum-to-one, and so you can think about them consistently across products. So, if I tell you that I've got two products here, one has a cross elasticity of 0.01 and one has a cross price elasticity of 0.03, I don't know, you don't know which one is the closer substitute, right, because you need to know market shares of those products in order to figure that out. Diversion allows us to do that. Okay.

Where I'm going to go today is talking about diversion as sort of a helpful compliment to things like merger simulation and also helpful additional variation in sources of data to help us estimate preferences in substitution patterns. So, the idea is that diversion gives us really nice variation on switching, on substitution directly as opposed to market shares and prices and characteristics. So, what I really want you to take away from my time up here today is the importance of leveraging that additional information for any of your analyses. Okay.

All right. So, the first part of the talk is going to give me ... There's so many numbers up here. I'm sorry. There's like a counter that's not counting, and there's a clock that switches on every once in a while for me. So, I'm evidently very easily distracted up here. So, the first part of my talk today is going to go through diversion, some results that Chris and I worked out and published in RAND last year that tells us what the empirical content or the empirical properties of diversion ratios are. That's going to be really helpful then for some new work that I want to talk about in the second half of the talk and which is the work with Paul where we're going to work on estimating preferences and substitution patterns.

I'm being a little careful to say preferences and substitution patterns as opposed to saying demand because I'm going to not actually work with price, okay? So, sort of preferences and substitution patterns when we can use only data on switching or data on second choice survey style data, okay? So, that's going to be a product space approach, if you will, to demand estimation, but before we can do that, we need the results of the RAND paper, okay?

So, the RAND paper works out the properties of these diversion ratios, the empirical properties, right? If you think back to the 2010 Horizontal Merger Guidelines, there was a lot of discussion about how we think about these demand ratios empirically, right? Do they come out of the firm's business materials? How do we measure them, or do we estimate a demand system? So, of course all of those answers were correct. There are many different ways of thinking about how we would measure a diversion ratio, and different measures of it is are going to give different estimates of diversion.

So, here we are. The diversion ratio is this, the fraction of these switchers, right? So, suppose let's think about diversion for now as raising the price. So, I raise the price of product J. Certain number of people are going to leave, okay? If you are not purchasing product J, you are treated ... Okay, I'm going to give the local average treatment effects version of this. The outcome is the fraction of those consumers who switch from J to K, and compliers are those who would've purchased at the lower price but no longer purchased at the higher price, okay? So, this allows us to write down diversion as a Wald estimator, and we can then ... I'm going to give you this representation of diversion, which is going to have two components to it. This is going to be the key to understanding the empirical content of these things, okay?

The first component I'm going to denote, DJK [inaudible 04:03:32]. Okay. This is sort of a structural parameter. This is the individual diversion ratio. My individual diversion ratio for cars, right, might be that my first choice is a Prius and my second choice is a Honda Accord, and my third choice is a Toyota Camry, okay? There's another customer, Bob. His diversion, his ranking, his individual diversion ranking might be something like his first choice is a Prius, his second choice is a Corolla, and his third choice is a Civic, okay? So, these are independent of anything about the intervention that leads us to leave the Prius. Okay?

So, the second component here is this weight, okay? The WI. This is going to depend on the intervention. So, the weight is which of those people I'm treating. So, if I increase the price of the Prius by $500, maybe I treat Bob but not me. if I increase it by $6,000, maybe both Bob and I get treated. So, the overall diversion that we're going to measure in the economy or for that particular intervention and as a result of a merger, say for example, is going to depend on which of those individuals get treated. So, that sort of distinction that we can think about those structural individual diversion ratios, sort of independently of the intervention, and then think of the weights as arising from those interventions, that's going to be key to

Julie Holland Mortimer:

... how we can utilize data on these things. So the MANN paper sort of establishes this and derives results for the case of discrete choice, sort of logit based estimators. So if you think about a plain IIA logit, diversion is just going to be, everybody's diversion is the same and it's all proportional to share, so that's the assumption sort of the equivalent in a late kind of framework that everybody has the same treatment effect. Homogeneous treatment effect. The weights that as that are associated with different types of interventions may also be derived, and so this table is from the Rand paper, so second choice data, I'm going to wait you based on the probability that you were going to buy product J in the first place. Okay? So S I J over one minus S I J. So this is just how likely of a buyer was I? How strong where my preferences for the Prius.

If the intervention instead of is a small price change, the people that I'm going to pick up initially are going to be people who are very sensitive to price. If it's a quality change, it's going to be the weights, how I'm going to treat people based on how sensitive they are too quality. And you can do the same thing for finite price changes and so on. Any characteristic that you want to think about. So now we've established those properties, I want to talk about the use of second choice data. So we're going to think about this first line in the table, if you will. Okay? So the second choice data, I'm going to use those weights, and here's the idea, Okay? So when do we want to think about using second choice data? So it can be a lot of cases where we might have different forms of second choice data. We might have an online company that's doing a lot of AB testing where I show different research results to different consumers.

I might be doing sort of conjoin analysis work and marketing. I might have merits or Camip style, second choice survey data. The automobile manufacturers often will run surveys like this. If you hadn't bought the Toyota Camry, what would you have bought? Rank ordered list is another example. I'm going to think about this as of a matrix of diversion. So think about this diversion matrix. So here's an example on the slide of a hypothetical, I sort of picked the numbers, the headings here, to evoke the idea of a T-mobile sprint merger. So suppose T-Mobile and Sprint are considering a merger, and T-Mobile has some business documents, some business research they've done that indicates that 30% of their consumers would switch to Verizon, 45% would switch to AT and T, 20% would switch to Sprint and 5% would switch somewhere else. And Sprint maybe says, okay, 30% of Verizon, 15% to At and T, 45% to T-Mobile, okay? And so the question is, suppose we had that data, that second choice data, and we had aggregate shares data and we had nothing else.

How far could we get filling in those sort of missing elements of that diversion ratio? How well could we understand preferences and substitution in this market with just that data? I'll make one last comment about this, which is that I've made these numbers sort of darker in lighter shades of blue, and I'm going to later show you some pictures where they're going to look sort of pixelated, and I know it's a little clunky, but darker shades of blue are always going to evoke higher rates of diversion, like a heat map. All right. So what do we typically do if we want to fill in this matrix? Well, we would have some parametric assumptions. So if I've got sort of nested logit model, I'm going to write down diversion, which is going to be a function of the shares and the within group shares. And so I'm going to actually use, if I've got six nests, I might have a rank six version of this matrix. If I've got a plain vanilla IIA logit, it's going to be just a rank one version of this matrix that's completely determined by shares.

And if we've gotten Random coefficients, we might be putting normal distributions on these taste parameters and estimating closer correlation and tastes for when we have high substitution between two products with similar X's You can think about this in the McFadden and Train context. One way to think about this is that we've done a lot of work, and many people in this room and online, have done a lot of work really helping us to get much more flexible on the distribution of tastes that are associated with X's. And we've probably done a less work on exploring what it means to have a sufficient set of X's. And because the way we're going to approach this in this new work with Paul, this is going to be a completely product space approach, so we're not going to have any X's. And so this is working to fill that out a little bit.

So the strategy here is directly come up with a low rank approximation to this diversion matrix. So we're going to be limiting it directly in product space. This has some advantages in the sense that it'll allow for sparsity and individual shares and substitution patterns. We've got the possibility of generating extreme patterns. If I'm at a vending machine and I pull out a Snickers bar, usually the portion of people that switch to the very top products is always greater than what a logit based model is going to predict. So we can get those kinds of extreme substitution patterns as well. So the type of low rank approximation work that we're going to do is the kind of work that's had a lot of success in the computer science literature. This is going to be, we're going to use singular value decomposition, to come up with low rank approximations of this.

And this is what, at least to my understanding, similar to the kinds of techniques that Facebook uses for image recognition, so sort of image compression. So this is a picture of Camille Jordan and all the way over on the left you see the original image, and then next to that is... And so that's sort of a matrix of pixels with rank 266. And then the next one over you can see is, after we've used the singular value decomposition, this is a matrix of rank 25 approximation. So this is the image compression idea and the approximation error. You can kind of see the shoulders of his jacket are not quite as well defined, and so we give up on some precision there. And so this is what our version is going to look like. We're going to have diversion. This is using automobile data that Charlie Murry kindly gave to us.

So this is in from his work with Paul Grieco and Ali Yurukoglu. And so all the way over on the left is the original data, so that's 173 different models. Actually this is just cars 66. So I think this has rank 66. And then in the middle we have what we're going to get out of our estimator, which has a rank of 13 and the error matrix next to it. Okay, so when is this useful to do? So first of all, when you don't have very good X's, it can be useful to do this or when the X's don't accurately capture substitution across products. It can also be really helpful when we're estimating substitution patterns and we only see sort of a subset of products. So I'm going to also give you an example soon from the context of vending machines, which Chris and I have done a lot of work on, where we were able to run field experiments.

So we were able to remove six products from the machines and see where people go. So we see a full column of diversion for each of those six products, but we don't see anything else. We could only experiment with six of those goods. So that would be similar to a merger case where you have really good data from the two merging parties, but you might not have it from the other firms in the market. It can also be helpful when, again, as a way of leveraging second choice data. So I said one thing that I want to make sure to impart with you today is the sort really a mission, sort of a challenge to practitioners and researchers to fully leverage data on diversion data from second choice, data from surveys that tell us directly about substitution patterns. And what we're doing in this work is a little bit different in the sense that's saying like, "Suppose we only have that data, how far can we get?" But they're obviously related issues and I think both very important ones. Okay? So here's how the model's going to work.

So we're going to use a discreet choice setup, and it's going to be semi parametric in the sense that we're going to be fully non-parametric on the indirect utility, the VIJ, okay? And we're going to try to estimate that by approximating with a finite mixture, with different weights. So if you remember in that Rand paper, okay? Well, I'll just give you the notation first and then tell you about the Rand paper again. So individualized share, I'm just going to denote as a vector, S I J, bold, S I J, and the aggregate shares now are going to depend on those individual shares, those individual probabilities and of purchase and the set of different types of consumers. So if I have two types of consumers in the market, it's just Julie and Bob. If I have three types of consumers in the market, it's Julie and Bob and Sue and four types, Julie, Bob, and Sue and Steve. So we're going to write down this matrix of second choice diversion, which comes directly out of that Rand paper.

So now we're going to have the number of consumer types is going to be capital I, and each type is going to be have weight pi I. So we're going to be estimating weights on these types, how prevalent are the Julie's and the Sue's and the Bob's in the market. And then we've got our individual diversion ratios, one for each type, and then we've got the weights, and those weights come right out of the Rand paper for second choice. But if you wanted to use a different type, fewer sort of changing quality, you could use a different weight. I've just rewritten that in the lines below in matrix notation and for the whole diversion matrix rather than just a single pair-wise notion of diversion. And so here's this diversion matrix, and here's the nice thing, is every individual diversion ratio is of rank one. So it's just the outer product of SI with itself. And so now we've got this unrestricted matrix of diversion ratios that would be of rank J plus one by J plus one, but now every individual type is going to be of rank one.

And so the rank of my compressed image version of diversion is just going to be given by the number of types. So I might have 66 car models. If I can approximate those patterns with 13 different consumer types, I'm going to have a rank 13 approximation. That's sort of the idea of this, using this low rank approximation to capture these patterns. I mentioned this earlier, but just to kind of fix ideas here, if it was an nested logit model, you would be imposing a rank of equal to the number of nests. If it was vanilla, plain vanilla IIA logit, that would be a rank one approximation and so on. I'm going to try to show you some pictures where I build that up. Alright. I'm going to just write down the estimator, I think on the next slide, and I think we've already covered this, we're going to just try to obtain an estimate for that matrix, that is of low rank approximation.

And so the estimator is going to be matching essentially two things. It's going to be doing our usual thing in the second term there, which is matching shares. So every time you see a script somewhere, that's the observed data. When you see the PI I, S I Js, that's our model equivalent, and so that second term is matching on shares. The first term is matching on diversion. So we've got to observe diversion. And then we've got our predicted diversion that comes out of those expressions in the Rand paper. There are a couple of additional things here. We can also penalize the number of types if we want to keep a more parsimonious model that will add a ridge style penalty and turn this into an elastic net estimator. Right now, so the CJ Tilda and the CJ's, are weights that we're using, that are proportional to shares.

This is sort of an ad hoc thing we're doing right now to make sure that we're matching well on shares. It's just different weighting, we don't have to do that sort of theoretically, but what I'm going to show you today does have those, and we're going to use cross-validation to select the number of types. So what we're going to be doing is we're going to be matching on diversion and shares, and we're going to continually increase the number of types and we're going to be computing median absolute deviation and root mean squared error to figure out when we've got enough types. And we're going to use out of sample validation for that to make sure, so that's our cross-validation process to avoid over-fitting to having too many types to do this. Again, we could estimate sparse models here. There's not these fat logit tails that are going to force us to smooth out diversion too much.

We can have pretty sparse models, it's a semi barometric logit model in the sense that those VIJs that I mentioned earlier are going to be completely non-parametric, and our probability of purchase is just the sort of logit form just for capturing those. Okay? There are related papers that have done similar things. So Fox, Kim, Ryan, and Byery have you bring the shares in and estimate the PI's. I can say, should say a lot more about all of these things, so apologies in advance. Davesh's work has bins for each of those different types. So he's got the PI I's because that comes out of the data for him and can estimate the probabilities of purchase. We're most similar to Green and Hencher which are estimating both the PI I's and the Beta I's. Okay? We're sort of estimating both of those things.

So there's three ways, and I think I have 10 minutes. Do I have 10 minutes? Is that right? Okay. So in 10 minutes we can cover three things. There's a Monte Carlo, this is all new stuff. So there's a Monte Carlo just to show you how the model performs and when we can control everything. There's an example using the automobile data that the Grieco, Murry and Ali Yurukoglu paper are using. And then there's a model with the vending data where we can show you against this model against parametric demand model and also show you a network graph. So I'd like to get to a couple of those things at the end. I'll go quickly through the Monte Carlo. So what we do here is we use the vending data that I'm going to show you later as the basis for generating our Monte Carlo, generating our fake data, so we generate fake sales and diversion data sort of motivated by this model. We do two things.

The first one, the truth that's sort of true data generating process is a Nested model. And the second example, the true data generating process, is a Random coefficients model that puts Random coefficients on nut content, salt and sugar. And then we compare those in sample to a number of miss specified parametric models. So suppose the truth is a nested logit but we estimate a Random coefficients on characteristics model instead and vice versa. And I'm going to just show you the out of sample predicted diversion ratios for our model. So I'm going to stack the deck a little bit against this current approach by showing you in sample predictions across the parametric models in an out of, against an out of sample prediction for this rank reduction approach. I'll show you median absolute deviation on root means squared air. And so here there's a picture rather than just a number because we're increasing the number of consumer types as we move across this graph.

So you can see that by the time we have about four simulated consumers, were doing really well. And the horizontal lines, of course, a nested LO model doesn't depend on, there's no type of number of consumers here. We're doing a couple of different things. We're comparing this here, the truth, the data generating process is a nested logit, and we're going to be doing better as we get closer to that. So at the very bottom here is Random coefficients on nests plus auxiliary moments with generated directly from the diversion statistics themselves. So just using them, those statistics with the parametric model directly so that we are bringing in the diversion information, but maintaining our parametric assumptions, so that Random coefficients on a nest is not very specified, and then when we add the diversion in, that model does quite well, but the semi barometric approach is doing quite well by the time we get to four consumers. Here's the same thing when the true data generating process is Random coefficient on characteristics and we're trying to fit it with a plain vanilla logit or a nested logit.

Again, by the time we get to four consumer types, we're doing quite well. So here's the application to the autos data. So this is a subset of the data from the Grieco, Murry and Yurukoglu paper. We're going to focus on one year of sales from 2015, and this is second choice survey data from merits where I think we actually have constructed these to be 173 products. Now we're going to predict unobserved second choice data again in product space with no characteristics. So this is what the merits data look like, again with that sort of heat map idea that higher levels of diversion have darker blue. This is the cars in the top left. In the middle are the trucks and some more consolidated products further down. I'm going to zoom in now to the cars. Okay, so that's here and it's probably, I can't read it from where I am, but you can see there's a Corolla, a Nissan Altima and so on.

So each row here is a different type of car with diversion to other substitutes. So now we're going to estimate our model on this. There's not a benchmark for the parametric model I had in the Monte Carlo, but you can sort of trace out the number of individuals and how we're doing, and here we're looking like we're doing quite well with 13 consumers. So we've got a sort of 66 by 66 matrix that we're able to approximate pretty well here with 13 types. So a rank 13 matrix in terms of some other things that you might care about. The percentage of the top 10 products that we fit, we're doing pretty well. We're getting about 66% of those. And the pairwise matches, we're getting about 74% right. And this is kind of what it looks like, just to give you a sense of how we can build this up as we add additional types.

So in the top left hand quadrant there's the raw diversion, and then with one type you see what kind of diversion matrix we can estimate. And then the lower left hand corner is two types. And then we've got our 13 type estimate of that diversion matrix in the bottom right hand corner. And then last few minutes, I want to talk a little bit about the vending data. So this is the same data that Chris and I worked with when we studied vertical contracts in vending. So there are 66 vending machines placed in downtown Chicago and we've got 35 to 40 snack products in each building. And this is the example I mentioned earlier where we can ask, we were able in a field experiment to remove six products. So we remove six products, one at a time, sometimes two at a time, so there were sort of eight different experiment arms there and we see where people go. So now this diversion matrix is going to look a little different.

It's not a survey from merits where we might have missed some people or we might have noisy estimates rather. The first two columns are Dorito nacho removal, Cheetos removal, and then Snickers removal, peanut M and M's removal and so on, and then the cookie ones. So now we need to fill out the rest of this matrix having only six columns of data. So we're going to go through the same median absolute deviation and root means squared error that I showed you before. And now I can estimate this against some parametric models, so I'm going to show this to you against a plain vanilla logit, a Random coefficients on salt, nuts and sugar and nested logit sort of chocolate candy bars nest and a salty snacks nest and so on. And so relative to the parametric models, we're doing quite well. What I'm going to show you now is at about two or three consumer types. So by the time we already get to two or three, we're doing pretty well.

And so again, just to build this up, so you know what this looks like, when we have two consumer types, when we have three consumer types, when we have four consumer types, we're getting more sort pixelated sort of more granular estimates of that diversion and the products are ordered, you'll see them actually, well on the next slide you can probably already guess. So these are the parametric models. So we've got the raw data on the top left corner, and then we've got our rank three model, and then on the bottom left hand corner is Random coefficients on characteristics and then Random coefficients on nests. And so the Random coefficient on characteristics, which has a rank of 10, you can see the characteristics that drive those substitution patterns. And in the nest you can see salty snacks is the first sort of block diagonal chunk there, and then chocolate candy and then non chocolate candy and so on. So you kind of see what those parametric models impose with the load error term.

And then the very last thing I want to show you in my last two minutes is we can think about these things because diversion always sums to one. We can think about this in a network sense. We can sort write down what the network graph looks like and that's what I'm going to show you next. And so this is the network graph for diversion where all the edges here have average diversion in both directions, greater than four and a half percent. And so lighter shades are larger share products. So if they're circles, they happen to be salty snacks, if they're squares, they happen to be chocolate candy bars. But that's just ex-post, right? None of that is going into the model. The model is estimating diversion and this is what's coming out of it. And so a couple of my favorite things that come out of this, way of looking at it is you see things like sun chips, and is sort of in the center, and then just to the left of sun chips and up a little bit is an unlabeled green one, which is planters peanuts.

And so it sort of gives you some insight, I think some unique insight, into which products here are likely to be central in a network sense and which may be providing more discipline, pricing discipline potentially in the market because they're positioned as close substitutes to both candy bars and salty snacks for example. And that's one of the things that I think is actually going to be really interesting to think about in this sort of new kind of approach is, can we help to segment different parts of the market by using this? We can of course add co variants or other things if we want to. And the last comment here on the extensions is sort of the optimal experimentations, which products would be good to remove, Which ones are we likely to learn the most from?

And this is something that I think is done quite commonly in the UK when they run second choice data and so on. So the question of which products are important for completing substitution patterns. So we think this is promising, we're kind of excited to make a push on this and learn what we can from second choice data and just encourage everyone to fully leverage data like that on substitution patterns whenever you can. So questions?

Speaker 7:

We have a moment or two for questions.

Ted Rosen:

Thanks, Ted Rosen, FTC. Does your method have a way to deal with sampling variation in the diversion ratios to scale those estimates to reflect how sure we are that the diversion ratio is out there?

Julie Holland Mortimer:

So that's something that we're kind of still thinking about. It's sort of very different in the case of the merits' data, for example, then in the vending data, sort of in the merits' data, there's more sampling error, but broader coverage. And I think we're still thinking about that. I think that's an interesting question.

Speaker 7:

Great. Well thank you very much. We will come back together at four o'clock for a keynote. Thank you.

Speaker 8:

Okay, it's time for our last keynote by Catherine Tucker. It's on data accuracy, digital exclusion, and inequality. Dr. Catherine Tucker is the Sloan Distinguished professor of management, MIT. She received her PhD in economics from Stanford and studies the economics of digital data. Thank you Catherine.

Catherine Tucker:

Wonderful. Okay. So it's very exciting to use Keynote. I'm so thrilled to be here. I also notice when I'm scheduled, and usually when you're the last speaker of the day, you're sort of battling against people's dreams of cocktails. But because this is the FTC, I've realized it's even worse, I'm actually going to be battling against people's need to get in their cars to beat DC traffic. So I want you to know that I'm very conscious of my responsibilities and I will aim to speak fast. I also know though that this means that at the end of the talk, no one is going to want to ask questions because they will just get murder looks from the entire audience. And so therefore, if you do have a question, I've managed to persuade poor Will to actually run around during the talk with a microphone so we can actually even be spontaneous.

We'll see how that goes. Or the other thing I should say is that I have consulted for every major tech company except for Apple. Now with that said, let us go on to the plan for the speech. Now what I'm going to be sharing with you today is free papers. And I would say that each of these free papers is very depressing. And so I thought, well, maybe I should start the sort of speech by saying something positive and mentioning that IT and digital data is a powerful force for good and also has an in times would use inequality. But then I actually remembered that we are all IO economists here and we love depressing. So why don't we just go straight to the depressing research and get rid of the, sort of ignore all the feel good stuff. Now I'm going to start off with my first paper. And this is a paper I'm very fond of and it is, I'm fond of it. I've got to say it's already published. It was published back in 2019.

I'm fond of it because it gets a lot of traction on Twitter, so I sort of feel part of that community. I'm also fond of it because what I love about it is that one of my co-authors on it, Tim, I should tell you who they are. Nico Neumann is an amazing professor and he's going to be my co-author for all the three papers that I present today. And if you haven't met him, he's absolutely brilliant. And he's in Australia at Melbourne Business School. Tim Whitfield though, I love too, because he's a co-author I've never ever met, I've never actually emailed him, spoken to him. I actually know very little about the mystery that is Tim Whitfield. You might say, "Well, how can he be your co-author?" The answer is he's actually going to be essential in the story because Tim Whitfield is a very powerful man and it is Tim Whitfield who's He's going to get us access to data, which is going to uncover an industry which usually we don't have much insights to.

So Tim Whitfield very much deserves to be there and he's going to get us such access that in some sense, you should also be expecting all the empirics to be presented in a very simple way because we're going to get access to data that no one else has had access to. So what we do in the first paper is we ask this question, which is, in our world of big data, in our world of machine learning, which is sometimes called artificial intelligence on PowerPoint, how good a job is it at doing at one of the big jobs that it's used for in marketing, which is providing segments of consumers to advertisers? And we think this is an important question because obviously I need tell a competition of foresty that there's a lot of talk about how valuable data is.

There's a lot of speculation about the role of machine learning in this world, and in some sense we're going to be unpacking the hood of a key element of the use of data and also the use of data too, which I would say also touches on a lot of people's privacy, stated privacy concerns, which is the use of data to profile you for marketing purposes. So we think it's an important question to study. What surprised me when I started off on this project, is that when it came to data in the economy, I had sort of been naivety assuming that the kind of data that advertisers was buying, were things like Catherine looked at orange minivans last Tuesday, something like very specific, something actionable like that. But instead, no, actually what I found out was a lot of the data that marketers are buying when they're trying to target individuals, and these are the profiles you're buying when you're targeting an individual

Catherine Tucker:

Individual outside of say Facebook or Google or Roku where the data's integrated. This is the kind of data I would be buying to target individuals with advertising on sites like ESPN or CNN or something like that.

So we are looking at that kind of data and it was just surprising to me as I say that actually what a lot of that data is it seems very Prozac. It's age, gender, household. In other words the kind of data that we've used for marketing and targeting for a very long time.

Now, this means because we want to try and understand how good a job this industry is doing in terms of providing this data, we're going to focus on age and gender at least in the first paper.

Now the question is of course, so this data is being bought from something called data brokers and there's a variety of data brokers in the USA ,and the FTC wrote a very good report on them, and we're going to be studying them from an empirical perspective.

What they do though is that remember the way that data brokers works is they really should be thought of as a data aggregator in that they piece together data about a cookie, or an individual. We tend to call them cookies in marketing. And they piece them together to try and work out what segment that cookie belongs in.

So the question is of course, well how safe an age and gender would a data aggregator which is getting these floods of data in say about your browsing behavior, what you've looked at. "How are they going to know your age and gender?" Well, let's sort of think about it as a prediction task because we tend to, when we think about machine learning. So I could use browsing data I guess to predict if someone was female, if perhaps they were browsing like I don't know a tampon webpage or something like that, that might tell me someone's female.

I might look if someone's browsing retirement homes, maybe that would tell me something about their age. But you can already see the problems. And for retirement homes I could be browsing them for myself, I could be browsing them for my parents, my grandparents.

In other words, I'm going to argue that actually using browsing behavior, people's digital footprints to predict age and gender is a really, really hard task. And I want you to have that in mind that we're looking at a hard task when I show you these results.

So what we're going to do in this paper is basically try and figure out how easy it is to get age and gender. And what we're going to do is... And this is all the power of [inaudible 04:43:05] Whitfield, is that we are basically going to get the data brokers to give us what their age and gender they think her cookie is. But we are actually going to know what age and gender that cookie is going to report themselves as being. Because we're going to be using this vast survey of around 30,000 people who've told us their age and gender.

Now we're going to go ahead and do that and then we're going to say, "Data brokers, given we already know what age and gender this person is, can you tell us if this is a man in the age bracket 25 to 34." And that's really what we're going to be looking at.

Now let me get you the results, and I promised you simplicity. And this is going to be the results we get from all the data brokers in our study. Now obviously they do remain nameless but you can think of them as representing all the major data brokers in the industry. These are multi-billion dollar firms. They are data brokers which are attached to major credit reporting bureaus. So they're big players.

Now we have this data for all the data brokers and then we're also going to have data on how well they do at the prediction task. And remember here I've asked them to say, "Is that cookie a man?" And what you're seeing in the last column is the number of times they accurately told me that cookie was a man. Oh go ahead. Will you get to run.

This is going to be our entertainment.

Pamela Meyerhofer:

Hi, I'm Pamela Meyerhofer from the Federal Trade Commission. Why does the number of cookies vary by data broker so widely?

Catherine Tucker:

Oh okay. Well, let's be spontaneous. Let's go to the second column. So what's going to happen is remember that we have 30,000 survey responses. And I'm going to go and say, "Oh for this particular cookie can you tell me if they're a man or not?" And in some sense, therefore we're going to see this as somewhat of a proxy for the data coverage of the data broker. And that's going to be important soon. But that's where that variation is coming from.

Now, when it comes to gender accuracy... I'm going back to my favorite, first column which is gender accuracy. And the reason it's one of my favorite columns is I want you to actually try and take a quick mean of that column. And remember this is the percentage of, times people, the data broker was actually able to tell me the gender correctly. And if you do a quick math you're going to get to just under 50%.

And if you think about it, getting gender right about 50% of the time is really quite amusing. You might say well how could they even be under 50%? And the answer is that sometimes they actually made us pay for the data and then they didn't return the results. So I counted that against them. Now the other thing... Oh we got another question, go ahead.

Zhenling Jiang:

Jenise Yang from Warton. I see we're going with this, but can someone argue that maybe what I care about is what type of things they're interested in? So even though I am young, I browse retirement homes and they think I'm old, but I do care about stuff in this category. So even though the prediction is wrong, but then you can predict my preferences right.

Catherine Tucker:

So let's talk about this because this is an important question which is sort of [inaudible 04:47:00]. They can't predict gender but do we really care? If you are a marketing professor, you're going to spent the last three decades telling all your students, "Don't use gender for targeting. It's very outdated." We're studying it because it's the main thing a lot of the marketers are using.

Now in case of what happens is even if they get gender right about 50% of the time, well maybe they're getting a hundred percent of right who wants to buy pink razors or whatever the product is you think gender based targeting is useful for. What we actually show though is even in the extreme cases where I think gender based targeting might make sense. The cost of this data is so costly. There's just no argument to use it rather than actually just taking... Well than just doing a scatter gun approach. And I'm going to show you some results which are more based on interests. Later on you're going to see some similar flaws.

So now we're getting to my exciting marriage of the columns though, which is the third and second column which is you'll also... The next math challenge for you in the afternoon is going to do a rough correlation between the number of cookies and how accurate they are. And you can do that sort of math of trying to see is there any correlation? I can tell you absolutely nothing.

And this was interesting for me just because we have a lot of theories about the role of data and how it might give competitive advantage or so on in prediction tasks. But it doesn't seem that coverage here is really helping you. And I'm going to argue here that this is maybe because it's just such a hard prediction to us that more data's not helping, but we just don't see that in the data. All right.

Now, actually the one thing I would just add is that... Oh no I'll go to the next part of the story. So we published this paper and I've got to be honest, it was hard to publish because it's a descripted result where we're showing an industry not doing very well. And so it's hard to publish in academia but it actually attracted a lot of attention from industry. And as a result of this, we actually teamed up with a very large firm that said, "Okay, so you've shown that sort of B to C data, traditional age and gender data's not very good, but what about business to business?" And if I'd had a prediction I'd have said, "Oh business to business it's bound to be better. If we got data on businesses, surely it must be better because we got more details. There's few of all of them and they must be easier to track all these things."

So what we did then was with the support of this large [inaudible 04:50:04] company, it was interesting that was in business to business sales. They were spending a lot of money buying data on something called IT decision makers. Very popular segment if you're targeting I guess in business of business sales. And what we did was we went out and again we repeated the methodology where we had ground truth about whether or not you're an IT decision makers, a little bit hard to pass. But the way we did it was we had surveys of people and said, "Do you actually seem to have anything to do with making a decision about IT?" And so the highest form was that you actually or ever in charge of actually paying for it or signing contracts a real what I would call decision maker. So we're going to use these ways of thinking about IT decision making and this is what we are going to find.

Now let me explain this graph because there's a lot on it. The lighter bars are going to be what we call the real decision makers, and the darker bars are going to be someone who looks a bit like a decision maker given their answers. And what we are going to show is that actually our ability to hit them that is target them directly. We're going to start off with what going to call our baseline case. And the baseline case is if I was to run an ad on a website which is devoted to business IT, how well do I do it actually targeting these people? And what you can see is there, I get them about 16% of the time. Okay so it's not bad. But then what we're going to do is we're going to compare that performance where you're basically targeting them through content with what happens if you buy data to try and target them elsewhere.

And the first kind of form of data we're going to look at is where the data brokers performing what we call probabilistic projections and it's all basically a sort of black box. We don't know why they're telling us these people are IT decision makers, but they're telling us there is a lot of AI and machine learning behind it. Then there's another type of data you can buy in this market which is where you buy lists. And this seems to me, it seems quite old fashioned I guess in 2022, but you can still buy lists of names of people and who are allegedly IT decision makers, and you can see the performance of that data.

Now what I want you to take from this is that given in both the probabilistic and the deterministic case, we were told that these people were definitely IT decision makers when we were buying the data. It doesn't look that awesome, does it really? In terms of how many people that we were targeting actually did seem to be IT decision makers given their survey responses. So that's the sort of first thing to take from it. I would say it's even less awesome than gender.

Now in the paper we spend a lot of time trying to think about why it is bad and what we show is that for the probabilistic that is the predictions of whether or not someone is an IT decision maker basically a lot seems to be going wrong with trying to match up their cookie with the cookie details that the data broker's giving us. That seems to be one of the reasons a lot of this is going wrong.

Now with the deterministic data, this is the long lists. Well, yes we have names, titles and emails but they just didn't react reality at all.

You can do a search on LinkedIn and none of it actually really matches. These are very outdated and [inaudible 04:54:15] data a lot of the time. All right, so from this I want you to conclude that it's bad in business to business too. In fact it may be even worse. So perhaps sort what you're trying to do is a little bit harder. The thing that really, really hurts me about the paper is that quite honestly rather than buying any of this data, you would do better if you targeted old men and that just hurts. So I put that in small type.

I guess the only good news is though that we know the age and gender data isn't any good anyway, so you can't do that either. But it certainly does speak to the sort of lacks of these methods that just basically vague, stereotyping might actually do better.

Okay, so hopefully so far I'd convince you that through data accuracy problems in both business to consumer and business to business market markets. And now I want to move this to being more so that's a descriptive result I think is interesting given what we think about data and the economy. But now I want to move it to being sort of more of a classic question in economics, which is sort of really actually going, I saw a show tell you this stuff is happening is to say well why is it happening? Why is it so bad? And what a distributional consequences might there be?

Now the amazing thing about this paper, and this is why I'm so lucky to work with Nico is that after our first paper we lost contact really with Tim Whitfield. That was sad. So we didn't have power anymore. And so what happened was that Nico went and actually worked at one of the large clients of data brokers and actually went out to get us a lot more data to help us delve a little bit deeper about why it was that we were documenting these sort of very poor results. And this is going to be in consumer markets on age and gender. And so we are now going to know a lot more about these people and that's going to be important for trying to understand what went wrong. Now the first result I want to show you, especially because we know we got IO, we might as well think about price. Is that typically when we think about markets we tend to think the high quality products would have high prices, and here the dependent variable is going to be the number of times that particular data broker seems to be getting data correct.

So we sort of changed the unit of analysis here and what you should see from this is that we have a price here which is how much you are paying for a thousand impressions worth of data. And you can see initially it looks good and it does look as though the data brokers are charging higher prices when they're more likely to be correct. The however happens though when I actually add controls for this, that effect basically goes away. And the reason it goes away is that data on genuine age is the cheapest kind of data. It also tends to be the most inaccurate kind of data on whether or not you're interested in sports, whether you're not interested in fitness, all of these kind of data, they're a bit more expensive and it's actually hard, easier to predict them, especially if you are trying to predict whether an Australian likes sports.

Most, I can tell you most of them do. So that's easy to get. And so as a consequence we can see here that once you start to control for the different things you're predicting, any relationship between price and accuracy goes away. So that's interesting and it was also a bit of a puzzle. So the next thing we did was, okay, well so if we're trying to think what's driving data accuracy or the lack of accuracy, maybe there's some kind of vendor characteristic which can help us understand it. And so what we then did was we ran a regression and you will have no hope of reading this slide, but let me tell you what it is. We ran a regression on how accurate the data aggregator was relative and we just had a fixed effect for each vendor and if you could see it, you would see there's absolutely no stars, there's absolutely nothing there.

I've tortured this data a lot to just try and see if there's anything which could help us give a supply side explanation of what's going wrong. The only thing I've ever found, and this is a very marginal result and I want to be clear, it is a result of data puling, but it is the case that if a data broker describes themselves as doing AI rather than machine learning that is weekly correlated with being less accurate. So that was the only thing I found. In other words, all the things you would math think would matter how big they are, whether they're connected to a credit bureau, how many employees they have, how long they've existed, all of the usual things you'd think might matter in terms of driving the quality of a firm just didn't seem to apply in this market.

Now though I want to get on to the promise of the title and so you should basically take from this. So where we got so far data is pretty inaccurate in both business consumer and consumer markets and we don't seem to have much going on in that supply side that we can sort help understand what's going on instead it seems to be pretty unsystematic and really related to the prediction task. Now what I went on to say, it's okay the but surely there must be consequences here of this inaccuracy and I thought it'd be interesting to try and study the consequences of inaccuracy for people from different social economic backgrounds. And so what I did here was, as I said, because of Nico's wonderfulness, we actually had a lot more data on this set of people and we knew about their income, we knew about whether they graduated from college, we knew whether they had known how sort of pretty standard measures of socioeconomic background.

And what you can see for this is we actually started to see a pattern emerge, which is that all the measures of socioeconomic prosperity or coming from an advantage socioeconomic background tended to mean that data about you is more likely to be correct. And here I've got a broad expanse of the kind of data which might be correct. It could be your interests, it could be your location, it could be your age, it could be your gender. But you can see here this was quite a striking result for us, which is that actually it does seem that rich people are profiled far more accurately and that was interesting.

Ofcourse we want to, that's a result, but we want to know the mechanism. The closest I can get to a mechanism is that there's sort of two explanations about what might be going on in terms of inaccuracy. And this is coming also from our work in the business to business market. One explanation is that if you are rich, maybe you just spend a lot more time online, maybe you just have a broader digital footprint and maybe that means there's more data about you should least more accurate predictions. Another potential reason for this is that if you are rich, maybe you've got a more stable digital identity. I'm looking at Steve Barry here. My guess is his email address has not changed for decades.

So some of you, maybe your email address has changed, that means you're probably young PhD student but for the older ones of us reached some kind of level of severity. Your email address doesn't change. You have a dedicated IP address and you tend to keep the cell phone forever. And as you can tell, we got, do we have a question? Are we getting excited? We got to make him run around. This is exciting. Anyway, this slide is trying to test for these two fairies and I'll explain it off the question.

Daniel Asma:

Well, there's another explanation why the richer people might be better profiled. If the marketers are trying to get the big spenders, then they're going to sort of [inaudible 05:03:03] regression's going to be driven by what predicts it for those then the poor people might not work.

Catherine Tucker:

So I love that and I was hoping for that because that was sort of more market based. And the problem I think comes from the fact that the kind of data that marketers are buying is not going to be very indicative of how rich you are. Woman, man, age. That's not telling you stuff. Sports interested, not telling you stuff.

We did have some data though on prediction of income. And if you think about it sort of mentally we would expect higher incomes to be predicted more accurately if that was to be correct. No, right? We just don't. I wish it was there. I wish it was there. You can tell I wish it was there, right? But it's a good question. Yeah, so no that's not true. Unfortunately that didn't turn out to be explanation. But what we did do was we managed to get a little bit of suggestion that it's actually just to do with the fact that richer people's identity is easier to combine together and is less fragmented.

And the reason I'm going to argue this is I'm going to divide up sort the different things you're predicting whether or not it's interests and with interests then presumably a lot more data will be more helpful for predicting whether I'm or not I like fitness. But we can see actually the correlation between socioeconomic background and the ability to predict interest is pretty weak. We're just sort of not seeing that.

On the other hand is things like your location, which are far easier to predict if for richer people it seems in our data. And so I'm going to give that... And you can also see this large coefficient known in your own house, which is sort of another measure of economic stability. And so I'm going to argue that though it's the best I can do of the day I ever hand. This is somewhat suggestive that the reason we may see these differences in the accuracy of data profiling for rich and poor people is coming from the fragmentation of identities.

Now I think a natural response to this is so goes so what? Isn't it actually really good that poor people that they have to be marketed too. That seems good actually as an outcome. But I'd also point out that though I'm looking at marketing data, these are the same kind of firms and the same kind of procedures which are being used for quite consequential decisions such as credit reporting, background checks, all of these kind of things, we might feel quite differently if these results go across to the accuracy of those.

All right, so now I've got my generic conclusions and my generic conclusions are just saying that data profiling generally seems to be working very poorly and it seems to be particularly working poorly for you if you come from a less privileged economic background.

But I actually have some provocative conclusions. And that's my privilege or my duty as a keynote. And my first provocative conclusion is going to be this, we talk a lot about privacy, I've studied privacy for a long time. But some of these results are making me wonder, do we have it the right way around? Maybe in some sense what I'm showing is that if you are rich, your privacy is more likely to be trampled on in terms of the accurate predictions being made about you in terms of your data. But if you're poor, maybe the bigger row is about data actually being vastly inaccurate about you. And so I want to introduce this term of data deserts into the privacy debate, and suggests that actually that's something else we should be worrying about the sort of algorithm make and data exclusion.

My next conclusion is that a lot of the work in AI or the economics of AI has been focusing on questions of algorithmic bias. But what I'm going to suggest is that though we get very, very caught up by whether or not the beta in a prediction exercise are really fair and representative and all these kind of things, as economists, we all know that that beta is a very small serve of the work and actually like 95% of the work in any empirical exercise is getting the data ready. And we don't yet have enough of sort a conversation, I think in a debate about how data errors, the problems of data matching the things that we all know as empiricists really could be distorting outcomes.

And with that, I'll just say the other thing I think it shows which is interesting, useful is that a lot of the debate about digital data and competition has been focused on if you have more data, do you have more sustained competitive advantage. But I think what my papers show is that actually we should be thinking a lot about the value of processing and the value of actually being able to construct an algorithm which can do any kind of accurate prediction. And so is this the bigger question when it comes to data? And if so, what does that mean for property rights over data of so much of the value is really coming from actually being able to build algorithm correctly. So with that one minute to spare, you are the gentleman at 439.

Speaker 9:

Yeah. So I've got a question. So you showed us that each individual data broker is not particularly accurate. But if I was to buy data from five data brokers, how much would that increase the accuracy relative to just one?

Catherine Tucker:

Ooh, that's really interesting. I'm not going to be able to give you a good answer. The closest I can get to this is we've tried to... What I can say is that their performance tends to be correlated in that they tend to do better for the same cookie and worse for the same cookie. So given in some sense you want them not to be correlated, my guess is you're going to do a bit less, you're not going to do so well by combining them. But I haven't actually done that specifically.

Speaker 9:

Okay.

Catherine Tucker:

There's a gentleman at 439, 30 seconds.

Daniel Asma:

Hello. Thank you. Hi I'm Daniel Asma. I'm at Spotify. So question, I'm interested to hear you expand on... So if we of accept the premise that data is more inaccurate for certain segments of the population, perhaps those that have lower income, how much of this do you think is something like a permanent result as opposed to something that you think is itself subject to change as the market evolves? And so I'm thinking about, so in the advertising industry, of course a lot of money flows through Google, Amazon, and Facebook and they have first party data as opposed to, it sounds like third party data. So I wonder how much of that, you think would change things.

Catherine Tucker:

So if you're in marketing you know we're moving to a post cookie world. And let me tell you what we've seen from the results. Basically I'm not that optimistic. And the reason I'm not that optimistic than from what we can tell a lot of this is coming from instability of email addresses and the fact that some names are harder to match than others. We've showed that this happens, for example with African-Americans, but names are harder to match. We show it also happens with Hispanic populations. Because there's more overlap of Hispanic names, and so the data algorithms can't process them so well. So if we go from cookies to a worlds where we're using things like phone numbers and names to actually do matching, I'm not as optimistic as I should be.

Speaker 10:

Great, thank you very much.

Tom Kotch:

Thank you to everyone who participated in today's sessions, both on the [inaudible 05:12:03] and on the floor here. We will be reconvening tomorrow at nine o'clock. For those of you who plan to be in person security will not be any easier. So please plan accordingly. Thank you.

Oh sorry. Apparently at 8:45 fives we'll be convening. So even more time for security to make sure you're here on time.