FEDERAL TRADE COMMISSION

THE TWELFTH ANNUAL

FEDERAL TRADE COMMISSION

MICROECONOMICS CONFERENCE

DAY 2

Friday, November 15, 2019

8:45 a.m.

Federal Trade Commission

Washington, D.C.

FEDERAL TRADE COMMISSION

I N D E X

WELCOME

MR. ROSENBAUM: Good morning, everyone. Good morning. Welcome to the second day of the Twelfth Annual FTC Microeconomics Conference. Before I introduce Professor Steve Berry to say a few words of introduction, just a couple of coffee-related announcements.

One is there should be coffee ready soon. It wasn't quite ready yet, so if you want to go out and get your coffee during the session once it's started, feel free to do so.

The second one is that yesterday there were a few coffee spills on the rug, which took some cleaning up later in the day. So just a quick favor, if your coffee does spill, just please let someone working for the conference know once it happens, and that way we can deal with it sooner rather than later on. The building management would appreciate it, so thank you.

And with that, it's my pleasure to introduce Professor Steve Berry, the Faculty Director of the Tobin Center, this year's cosponsor for the Microeconomics Conference.

(Applause.)

MR. BERRY: So Ted asked if I was going to say hello and welcome to all the early risers this morning, and particularly the ones that made it through the security line, which is an impressive thing. We are cosponsor. I should make it clear that really, you know, 99.9 percent of the credit for this conference goes to the FTC, to the staff, to the economists that help organize it, to the scientific committee, to the presenters and the discussants.

So -- but when Ted called, I was super happy to become a cosponsor. At one level, you could say it's a very sort of simple transaction that we get a little tiny bit of advertising for our brand new policy center at Yale, and in return, we get some sandwiches and a little bit of beer at the end of the day.

But that, I think, is not really the transaction that either one of us was interested in, which is really to try to build academic ties that run deep and are serious. I think this cosponsorship recognizes that between people in academia who are serious about policy and policymakers who are serious about getting their research into the policy agenda.

So I was going to take just two minutes maybe to tell you about -- a little bit about our new center. I come from a department which had two great centers of research -- one focused on methodology, one focused on sort of international matters. And it's probably always been true that economists should be contributing to the domestic economic policy debate with nonpartisan and evidence-based research, but this seems like maybe a particularly good time to try to get people to focus on actual evidence and to see if there's anyone we can get out of their corner.

So our idea was that we would be really based on economic research, that it would be nonpartisan, as the policy center people say, that we would try to at all times focus on evidence-based policy rather than on policy-based evidence. And I have to say, if you look around the country, I mean, you get a mix of kind of university policy centers, some of them definitely are located, and I think this is fine to have some diversity in this way, some of them are located pretty firmly in a sort of policy -- point in the policy space, right, and have a tendency to organize their discussion around that point in the policy space. And I hope that as I think the people at this conference do that we can avoid that, that we can actually let the research go where it does.

One kind of center that I think has been super successful in focusing on evidence-based policy are these centers that focus on kind of strict policy evaluation, that you see some policy, it's a pre-K program, it's a teacher training program, it's a particular way of giving income support. You see the policy, it applies to lots of individuals, maybe you need an instrument or not, depending on the degree of randomization, and then you see whether the policy worked.

And I think that's a great style of research, and there are people at Yale who are going to want to do that style of research, but if you look at my colleagues and you think the research is going to come out of the faculty, I think there's really a lot of other kinds of evidence-based research and research-based policy that we should be engaging in.

And, again, I think I'd come back to what's going on at this conference. So our President, for example, he's always saying, well -- he points to us and he says, you know, a direction for the future is big data, data -- you know, data-driven policy analysis. And that's fine, but under my breath, I always say, you know, and also theory because -- and I think we saw this yesterday with the talk we saw yesterday on deception, which is that theorists in many cases are really required to give us the vocabulary and the framework even before we can start answering questions.

So one of the first questions where someone from the policy community, someone interested in digital markets, how should we regulate them, how do we think about markups, came to us and he said, you know, I go to the Washington policy community and they say, Google Maps is free. I mean, you think there’s an antitrust concern? The price is zero, it seems like. It seems like you should just go away.

And this person say, but I think there’s a markup there. There’s like a trade; there’s an exchange of data for something of value. There could be a markup. And my theory colleague raised his hand, and he says, well, you know, that’s what I do. I have papers that define the data markup in the presence of a data externality that they’re learning about you, not just from you but also from all the people around you. And then this gives a framework where we can start to define -- I was, like, well, how do you -- how would you -- you can’t think of measuring the data markup until you’ve defined the data markup, right?

And I really think that one thing we can do, similar to this conference, is sort of to welcome theorists into the discussion of policy and how do we build a research agenda around policy, and I think this conference does a good job -- does a good job of that.

Everybody knows, though, what the President is talking about, which is that I look at my younger colleagues now who are often combining data sets from four different sources, they're all confidential, and they're basically doing the same thing that Google does, right, is they're learning about you and about the world by, you know, combining data from credit bureaus and address data and tax data and all kinds of things that are going on like that.

And I think that gives us the ability to, you know, in the first place, just describe the world and just tell us in a more detailed and more convincing way what's going on, and so I think that's another kind of research that people often don't stop and I think actually spend quite as much time on, which is just to frankly say you're describing the world.

And, you know, you see the paper and they say this is merely a description, and then they go to the table, and they say, and in this pure descriptive paper, we see that the effect of Variable 2 on Y is, you know, this, and I think if we can encourage people a little bit to take those big data sets and pause for a minute and not jump immediately to causal effects or whatever they are and tell us the way the world is that that would be a -- that would be a super useful thing to do.

And then, finally, I think something that I hope we're set up to do and to encourage something else that you see at this conference, which is counterfactual policy analysis, analysis of a policy which has perhaps not happened yet, which is obviously different than going out and using the pure variation caused by the policy quasi-randomization to learn about policy.

And, of course, once you think that, you realize that actually many even ex post policy evaluations are actually counterfactual analysis, right, that you're actually trying to recreate the world that would have been if the policy had not been undertaken, right? So, you know, you can ask what's the difference between, say, a prospective merger analysis, where you're very much trying to predict the world that will happen if the merger occurs, and a retrospective merger analysis, which is you're trying to predict the world that would have occurred if the merger hadn't been allowed, right?

It’s still very much a counterfactual analysis, and particularly in this case where the policy is complicated and involves an equilibrium response in either way, that it’s not just a stimulus to a person but rather an intervention in an equilibrium market. Actually, they’re both very much counterfactual policy analysis.

Now, it’s true, I think, that as in, you know, the study of the pre-K program, that the existence of the merger creates some variation in the data that’s not present when you haven’t seen the merger, right? And that source of variation can be very important in performing the policy counterfactual, but still, it’s a kind of -- it’s a kind of policy counterfactual analysis.

I think that you see all of those strengths here among the economists that are here -- you know, the ability to collect the big data, an openness to theory, helping us figure out what the world should do, a willingness to do counterfactual policy analysis of policies that we have seen and policies that we have not seen in an equilibrium context.

I hope that we can build on this partnership and maybe even build it out a little broader, to a broader set of economic questions. I think, for example, the tax policy community in DC is pretty sophisticated, but there are parts of transportation analysis, parts of environmental policy analysis, where they're actually doing incredibly complicated counterfactual analysis, you know, what would the urban residential and transportation patterns look like with or without a major improvement in the -- in a public transportation network is a massive policy counterfactual. The policy counterfactual of what happens under different environmental regulatory policies is a massive equilibrium policy counterfactual.

And there are communities of people trying -- very sincerely trying to do this in DC and elsewhere with very little input from the academic community. I talked to someone in the transportation world who was talking about trying to maintain their 1989 FORTRAN program for the cost-benefit analysis of a highway that no one knows what it does anymore, and some guy finally volunteered just to make sure the thing cranks and doesn't die, where people have really not had the benefit of this kind of back-and-forth analysis that goes on in this room.

But for today, we're all here, and it's so happy to see everybody on the same page, I think, looking for answers that can come out of the research and that we’re not precommitting to and being open to a methodological diversity that encompasses theory and data and modeling and counterfactuals, and I’m looking forward to the rest of the day.

(Applause.)

PAPER SESSION

MR. KOCH: So we will now move on to the paper session for this morning. The paper session was chosen by the scientific committee member Mark Schankerman. The first paper will be presented by the name on my phone, Yizhou Jin from University of California at Berkeley, presenting a paper joint with Shoshana Vasserman, and it will be discussed at the end by Allan Collard-Wexler of Duke University.

MR. JIN: Okay. So my name is Yizhou Jin. Thank you very much for coming. Thank you very much for the committee and especially for Mark for having us. This work is joint with Shosh Vasserman at Stanford.

Okay, so my research agenda in general looks at how and the process of which data has become -- a certain type of data has become available to certain type of firms, right? Especially markets with information and search friction, and further, how does this change in sort of information structure of the market really impact pricing, some market structure, and consumer welfare.

And in this paper, we're going to focus on a very -- what has become a very prevalent way in which consumer data are made available to firms, which is through direct transactions in which the firm sort of incentivizes the consumer to voluntarily reveal something about themselves, but on the other hand they also keep the collected data as proprietary. Okay, so this growing problem has mostly been attributed to two factors -- the advance in information technology and the strengthening of privacy standards. The latter really makes sort of voluntarity and consent essential to this process.

So we're going to go back to these two factors in our analysis, but let me first talk about an example, which is exactly what we're studying in this paper, which is the introduction of monitoring programs in U.S. auto insurance. So in this program, the insurer will invite new customers to voluntarily plug a very simple device in their car that tracks and reports how they drive for about six months. And in exchange, the insurer will use the data to better sort of assess accident risk and adjust consumers' insurance premium going forward.

Now, there are other examples, like in the North American life insurer, John Hancock, has a large program called Vitality that tracks people's daily health-related behavior in exchange for discounted life insurance. The Chinese tech company Alibaba has a proprietary credit score that’s linked to various price -- various -- the prices that you’re going to get -- you’re going to get on various rental and sort of borrowing services. And the way for you to improve that score is by giving Alibaba more data, like setting up your direct deposit or pay utility bills.

Now, outside of this sort of insurance landing selection market context, we also see, for example, Uber offering a credit card to its consumer, and it pays them much more to use this card intensively than what they’re going to make back on transaction fees.

Now, there are some other reasons for why they do this, but according to their term and services, one of the main reason, rationale, could be that they can link this individual transaction data back to their main business in ride-sharing and in food delivery.

So back to our main application. In this $260 billion industry in 2017, which is U.S. auto insurance, let’s think a little bit about what is the profit and welfare impact of introducing this monitoring program. okay? To answer that question, we acquired a proprietary data from a major U.S. auto insurer that runs one of such program, and, in fact, has introduced in a staggered fashion across states during our research window. And to further understand the competition in the industry, to Steve’s point, we match this data set to competitors’ price menu based on information from state regulatory filings.

So our empirical strategy, you can think of it as a two-step approach. First, we tried to think about how useful is this monitoring technology. And given that this is what we’re working with, we’re going to see how -- we’re going to ask how much information is really revealed in equilibrium.

So for the first part, we’re going to give you some reduced-form evidence that quantify the monitoring’s ability to both incentivize drivers to become safer as well as to reveal persistent differences in terms of the accident risk of different drivers, right? And -- or in other words, to improve the firm’s risk rating.

And for the second part, on the demand side, we need to estimate some structural parameters that governs how consumers’ information or monitoring opt-in choice is going to correlate with their insurance choice in the product market, as well as the cost to insure them because we’re in a selection market context.

And on the supply side, the key is to realize that the firm’s information set is not endogenous to the prices that they set for a monitoring program, if I attract more people into the program and I got to see more about this person and potentially gain a competitive advantage over my competitors, right? So we’re going to use a two-period two-product model to characterize pricing in counterfactual equilibria.

So zoom out a little bit. Conceptually, we’re essentially trying to first think about what is the sort of degree to which this information technology can address information problems in a market. And then, secondly, because stricter privacy standard says you must now purchase consumer data, with their consent, we’re going to try to use existing IO tools to not only quantify the price and quantity in the product or insurance market but also those things for the consumer information that monitoring is trying to pick up.

So we’re going to run two main counterfactuals. One is a no-monitoring counterfactual that help us infer what is the impact of introducing monitoring; and the other is sort of given that we observe the resource cause of this company running monitoring, we're going to see what's the optimal pricing that the firm should have charged, as well as on top of that what if as some of the regulatory proposals are saying we mandate that this proprietary set of monitoring data be shared with every other firm in the industry and, therefore, sort of eliminate proprietary data. Okay?

So I'm going to start with some simple background information. Now, suppose someone comes to the firm at Time 0. You need to make a coverage choice right away, and then each period lasts for six months, at the end of which, you need to think whether I stay with the firm or not. And the firm will give you a renewal offer to facilitate that choice at the end of month five.

Now, suppose I got into an accident. I will call to file the claim right away, and then depending on the claim type, pay something out of pocket, and then a claim adjuster will come here to evaluate the situation and give me the right amount of reimbursement. But very importantly, as soon as I call to file the claim, this information becomes public in the entire industry. Now, it goes into a shared data base. So my renewal offer, not only from my firm, but from every other firm, will reflect the fact that I have gotten a claim and, therefore, may be a more risky driver.

So for the first period, we're going to see observable characteristics of the driver. The quotes that they receive on liability limits, which are mandatory by states, varies between $30- to $500,000. It means in event that you are sued for liability, the company will cover you up to that amount, and your out-of-pocket starts thereafter. And then because prices are regulated and we have all of the observables that goes into pricing, we can match our micro data with competitors' price menu to see what are the competitor quotes that you would have gotten had you went to another firm. We also see the coverage choice and the premium that they paid for that coverage.

So at the end of each period, we're going to see claim realization. The average person have about one claims per ten years, and we also see how much your renewal quote changed compared to your current period prices, as well as whether you stayed with the firm or not.

Now, suppose you participate and after monitoring is introduced, you need to make an opt-in choice together with coverage choice, and if you do opt in, you have to accumulate about 100 to 150 days of monitored driving before renewal offer. And if you do, the company will include the appropriate amount of discount or potentially even a surcharge together with the renewal offer.

Now, in the first period, the company is going to tell you what are the monitored behavior, like your mileage, like late-night driving, harsh braking, and the monitoring duration only happens in the first period for new customer. And the company will sometimes give you an opt-in discount for the first period premium just for opting in.

And then the renewal discount, the company will also give you a range of the renewal discount, something like maybe 40 percent discount to like a 10 percent surcharge. That will be applied for all periods after monitoring ends, which is starting from the second period to forever.

Okay, during monitoring, they’re going to give you some real-time feedback, things that are -- suppose you slam on the brake, you know, your phone -- sorry, your device might beep or something like that.

At the end of sort of -- when we observe the renewal offer, we see a numerical score and the corresponding discount that’s given to each monitored driver. Very important thing to realize is that this is proprietary data, and we actually verified this information with filings and did not just assume that.

So in the interest of time, I'm going to really quickly go over our reduced-form evidence. It's essentially saying that monitoring is useful in two ways. One, drivers really become a lot safer, 30 percent safer, when they're being monitored compared to when they're not. This alludes to an incentive effect that says drivers can exert effort to signal that I'm a better driver going forward with the firm. And we are getting that with some -- within driver comparison in terms of claim realization during versus after these people that opt in sort of -- during and after when they're actually monitored.

But despite this behavioral change, we still see that the firm getting this monitoring signal at the end of -- in period one is going to be able to use it to better price a different driver's risk significantly. For example, like receiving a score of 170 deviation above the mean is going to be associated with 29 percent higher claim count in the subsequent unmonitored period, even if you condition on everything else that the firm would have observed otherwise, which includes the claim realization in the first period. Okay?

Now that we sort of have a sense about what this technology does, it’s important for us to have a model to -- a demand model to think about how people opt in and how this opt-in choice correlates with their insurance choices and the cost to insure them.

So I’m going to give you an overview of what this model is and what are the key parameters. So first, we need a claim model -- sorry, cost model that basically says what are the risk of any insurer is trying to take on and how risky are the drivers. In particular, we’re going to try to explain the arrival of claims, C.

And then we need a model for our monitoring technology that really tell us what is the information content that is contained in this monitoring signal, right, that influence the attractiveness of this program, both to the firm and to the consumers. And then both of these are going to go into our choice model, which includes product choices, whether you choose my firm or not, and what kind of insurance liability limits, insurance coverage you choose, and we’re going to use Y to denote that. But more importantly, we also want to be able to model the information choice, which is whether they participate in monitoring or not.

So for the first cost model, we’re just going to say that everyone has a latent risk type that partially depends on sort of your characteristic, like how old you are, and then but conditioned on that, there’s still some sort of unobserved heterogeneity that’s denoted by sigma-lambda here. And very simple way to capture this incentive effect that we just discovered is to just say that the consumer can change this lambda by some fixed amount, theta, when they’re being monitored compared to when they’re not.

And then for the monitoring technology, we’re just going to model this monitoring score, S, as an informative signal of this person’s underlying risk at hat-lambda, so with some precision sigma-S. So if sigma-S is zero, then you know that they’re observing this score S is equivalent to observing lambda, given that the slow parameter is nonzero.

And then for the choices, I think our product choices are modeled similarly to the literature in the sense that sort of your insurance coverage is determined based on how risky you are, as well as your risk preference -- risk aversion term, gamma, but there’s also pretty big inertia to switch firms that is pretty empirically sort of proven, so we’re going to have that term, eta, there, that sort of prevent people from switching between firms easily. And for the information choice, we can use existing parameters that we already have to try to model this financial risk and rewards very well.

Firstly, you drive better when you are monitored. So sort of you have some risk reduction, less likely to pay out of pocket, but on the other hand, you also receive a noisy sort of renewal discount based on monitoring, right, that depends on how good of a driver you really are, as well as how good of a signal that monitoring sort of score is, right?

But on top of that, just because it makes sense for you financially to participate doesn’t mean you actually do. So an important part of the paper is also this unobserved disutility that we need to specify that push people of -- even of the sort -- of the same observable group to differentially push people sort of away from monitoring.

So I only have ten minutes, so it pains me to have to sort of skip some of this, but I think in order for -- to really understand the structure of our paper, think about us being -- doing -- trying to do two things. We are essentially specifying a simple --

and introducing some theory to specifying sort of like a simple parsimonious model to achieve two things. One is that we have a giant choice base. Every firm offers a bunch of coverages, and after you have monitoring, you can choose monitoring with any sort of insurance coverage, right? So we're essentially collapsing that choice base based on the financial characteristics of sort of what is being covered when you get into an accident.

And secondly is there are two main sources of risk here. Suppose I'm a five -- like there's 5 percent chance that I may get into an accident, then whether -- there's a lot of uncertainty first in terms of the accident risk, which is to say that is this 5 percent going to realize this period, right? I want to cover that.

Another source of risk is reclassification risk, which is to say that because we have this information asymmetry problem, just because I'm 5 percent doesn't mean that the firm is going to think I'm 5 percent, right? So if I got into a claim or if I got a really shitty -- sorry -- a really bad monitoring score, then I may, like, you know, be punished dynamically in -- sort of in the future in the form of a higher premium. So essentially that's what our sort of structural model is trying to consistently account for.

Okay, essentially, what our model is going to be able to do is that this is empirical distribution of the monitoring score in the data. We achieve a pretty good fit, but you can also infer what are the people -- have everyone participate in monitoring what's the alternative counterfactual distribution that you're going to see, which is this sort of orange dotted line.

So you can see this clear advantageous selection here into monitoring, which is reflected in this disutility of monitoring term that we see. So not only is the mean of this term very high at $93, which means that the average person needs to expect more than this to participate, this is also higher for risk here, people, which means that even conditional on objectively what you're going to get from monitoring, safer drivers are still more likely to participate, okay? So it's important that this term be very flexible.

Now we can run some counterfactuals. For the base -- for the first one, we're going to run a no-monitoring counterfactual, which is we are going to hold baseline prices fixed, so introducing monitoring is not going to change your baseline, unmonitored price. We verify this with an event study. And then we know the resource cost of monitoring and we set it at $35.

So this is the change in welfare when you minus the -- sort of subtract the no-monitoring sort of regime from the current regime that we observe. The gray bar says the total surplus goes up by $13 or 1.5 percent of premium per person in our data set per year. And then on the left side is breaking down into an increase in consumer surplus, increase in firm profit, and a decrease in competitor profit.

So -- but perhaps more interestingly, if we get rid of the incentive effect -- remember, drivers drive 30 percent better when they are being monitored, right -- so that's a big source of welfare or surplus for us, but if we get rid of that, drivers are no safer when they're being monitored compared to when they're not. This is what you're really going to see.

So you can see a big part of it, at least we are -- this is a one-year horizon -- a big part of the short-term surplus that we get is coming from the fact that consumers behave differently, but another point that you can see, because taking away the incentive, we're left with the allocative effect efficiency improvement, right? So you can see that sort of the overall profitability of this market actually drops, which, you know, going back to the classic Rothschild/Stiglitz-type of insurance cream-skimming type of paper, which says that in the presence of information asymmetry, sort of competing insurers, trying to poach, like, better and better drivers without knowing that they are better and better, can only do so by offering less and less insurance coverage and, therefore, unravel the market.

But what we are showing here is that when they can compete based on information, they can sort of really attract good drivers with lower prices and, therefore, by transferring some of this surplus to the good consumers, push the market sort of towards a sort of perfect competition, perfect information, first-pass benchmark.

So, okay, good, now on to the pricing and equilibrium. So we need to specify a model to account for how the firms price this monitoring program, and we want to do so in a simple fashion. So we're going to use a -- first specify a two-period two-product firm profit model -- function. Two-period is because we want to cover pre- and post-information revelation. You don't just see this person is good in the first period when you try to elicit information, right? And two-product is because when you introduce monitoring in a voluntary fashion, sort of your monitored pool is going to cream-skim your unmonitored pool.

And for the firm's action, we're going to focus on three types of price adjustments that are specifically related to how the firm -- how the monitoring program can change the firm's information set. So in the first period, you know, the firm does not observe anything about this driver yet, so the only thing they can do is to either surcharge the unmonitored pool to sort of nudge you into monitoring or to discount the monitored pool to encourage you to participate.

But in the second period, once I see that you are 50 percent better than what I thought you would be, right, last period, there's a question of how much of that rent do I share back to you, like do I give you back 30 percent or do I give you back 20 percent, right, because you're already at my firm, so statically I probably don't really want to give you a lot of rent. Like even if you're 50 percent better, I might be pretty confident that you're -- even if I give you 10 percent back you are still going to stay with me, right?

But then dynamically, if you think about it from an ex ante perspective, sharing too little rent also will decrease the attractiveness of this program to begin with. So, okay, with this pricing model, we're going to run two counterfactuals. One is that we observe the cost of monitoring, so holding competitor price, we can always do optimum pricing for this monitoring program. How can you get the most amount of information to make the highest amount of profit?

And two is suppose we introduce this data-sharing regulation that eliminates proprietary data, saying you have to share this with other firms, what would you -- what's going to happen to the market? So, here, we're going to assume competitors have symmetric belief and profit function as the firm, and the action, we're going to only focus on one action, which is ex post to monitoring, they can set an alternative rent-sharing regime.

Remember the sort of 50 percent, how much do I share back that 50 percent? They can -- they can offer an alternative rent-sharing regime to poach really good drivers away, right? We really want this poaching sort of incentive to drive home the fact that monitoring now becomes a public good.

So I'm going to present the result in this table. You can see the first four rows are profit and welfare and surplus. The middle row is the monitoring market share. Think 15 percent of people opt into monitoring, but then we need to simulate an entire market out of which the firm only have a 20 percent market share, so the overall unconditional monitoring market share is only 3 percent in the data. So the pricing we're going to focus on unmonitored surcharge, opt-in discount as we talked about. And in the second period, there's a rent-sharing regime that the firm and potentially the competitor can set. We're going to benchmark that to one in the data.

So in the optimal pricing regime, the first thing I want you to focus on is that the unmonitored surcharge is only 2.7 percent, which is to say that when you try to coerce people into monitoring, not only do you push them into monitoring, but you also -- sorry, nudge them into monitoring, but you also push them away to other firms, right? Because auto insurance is mandatory, so the only -- like, the price competition is the only force that limit how much that can -- how much surcharge the firm can do.

So this is to say that price competition really does limit the ability of firms coercing people into revealing their information, which is not the case with Google and Facebook. Like post-GDPR, they really achieved a much higher consumer consent rate -- data consent rate than their competitors, and which is potentially not only because they have really good service but because their market power allows them -- market power in the product market allows them to contingent service among data consent in some cases.

But, instead, what this firm should do is sort of it really should offer a lot higher of an opt-in discount and also share less rent -- 80 percent of the rent -- in the second period, which drives home this invest and harvest dynamic that's pretty common in a lot of the ex post moral hazard -- sorry, ex post market power situation like, you know, like a network effect.

Okay, now, if we on top of that introduce data sharing regulation, you can see that the competitor offers a lot more rent back to the monitored drivers, which force the firm to also share more rent ex post, but this also decrease their incentive to offer opt-in discount in the first period, which drives down monitoring market share overall compared to the sort of previous equilibrium without this regulation.

So even though the firm is taking less share of the rent from monitoring, right, there is just much less rent to share in the first place, okay? This really goes back to a point first made by Richard Posner's 1979 essay, which says when data collection is socially valuable we should be careful about firms' property right to that data to protect their sort of incentive to produce that data in the first place.

So to summarize, drivers respond to financial incentives and become a lot safer. We got a very large incentive effect. Two, we find a strong -- there's strong advantageous selection into who reveals their information, however, not a lot of information is actually revealed both because we see large demand friction among consumers and because there's a lot of price competition that limits how much the firm can do to coerce people into revealing their information.

And given that -- given this competition, this sort of market structure, in the insurance market, we see that insurers property right, in this case, their property right to the monitoring data really strongly influence their effort to elicit data through pricing.

Okay, I'm going to leave you with this slide in which sort of I think from a policy perspective we can see that data regulation in insurance or the, like, broader privacy standard should really depend on the social value of the data collected, as well as demand and supply primitives in the product market, which says that sort of potentially requiring the disclosure of price or quantity of facts associated with certain data could be better than outright ban or full transparency.

From a research perspective, we also show you that information structure becomes an equilibrium object, just like market structure. So we shouldn't be sort of regressing other equilibrium outcomes on the amount of information in the market.

Okay, thank you.

(Applause.)

MR. COLLARD-WEXLER: So thank, Yizhou, for a very provocative paper to read. I enjoyed reading it quite a bit. So let me get into -- I think there's -- in the antitrust community and policy community more generally, there's an increasing thought about the market power effects of data. And you see that in a bunch of different places, for instance, Amazon Basics using all of Amazon's data to target exactly, you know, which product sector segments are going to come into, could somebody replace Google advertising, given that they have trackers across the entire internet that other firms have a lot of difficulty replicating, they just have a data advantage?

And so I think there's a thought that we need to think hard about the market power implications of data. And the insurance markets -- and I'm thinking here specifically things like life insurance or auto insurance -- these insurance markets have always been about what are the competitive advantages of data. They have collected data for a long time, so if you get a life insurance policy, they'll collect medical records, vitals, what you do, and so on, and this has existed for a long time. You know, life insurance companies have collected data forever, ever since, say, the 1850s when a large part of our capital stock was insured this way.

And I think what they're doing in this paper is saying what's the -- what are the -- what's the effect of data collection on equilibrium in these markets. So I think it's useful to separate this paper into two pieces. So there's one that's, I think, really like a treatment effect of the monitoring program, and then there's another one that's what is in equilibrium the effect of private data collection that gives one firm more information.

And so I'm going to give comments on one then the other, and, unsurprisingly, I'm going to suggest that these probably will be split into two papers at some point, so let me do that.

Okay. So the monitoring program can have effects in a lot of different ways. So the authors are very clear. The first effect is you just select better drivers into the monitoring program, and that might be about incentives or just which people want to sign up for other reasons, for nonpecuniary reasons, period. Then, you know, even among the kind of treatment effect of this monitoring program, it could be about financial incentives.

There's all this nudging going on, telling you when you're driving poorly, so it might not even be anything about economic calculation. It could just be the pure organization of the program, and then I think what's even harder for me to understand is what do people who are being monitored think the program is about because somebody's putting this device in your car and it's sending you all sorts of information on what you're doing, and so do I have correct beliefs about what is the effect of driving badly or not.

And I think with these very new programs that are very novel, the treatment effect you're getting from this first introduction might be very different than what if we had this device in the market for ten years where everybody kind of got used to it, a little bit like lane detection on your car. You know, the first time it beeps at you, you respond immediately, and then, like, three months later you start ignoring it. There's a real -- there's a real question of what are the behavioral effects of this device that might be outside of strictly financial considerations that are real.

So the one thing I think I'd be very careful about is when we're thinking about the treatment effect of this monitoring, you know, how much we think is coming directly from financial incentive effects and what's coming from other parts of the design of this device. The design's interesting. It just might not need the rest of the economic model to be analyzed kind of persuasively.

Okay. And then there's some -- some small comments like you need to restrict the sample to people who stay with the insurance company before and after. There's a whole bunch of drivers that kind of stop using the device halfway through. There's a comment in the paper which I have to ask you that, like, 10 to 20 percent finish the monitoring, and so there's a whole bunch of attrition that's a little bit complicated to understand that I think would be just useful to highlight. And there's no way you're going to put this into the model because it's just too complicated, but we'd like to know exactly how this data monitoring is kind of affecting behavior even if we can't put it into the model by itself.

And then, you know, other questions like what do you do with multiple drivers? It would be super nice if you had people that accepted the device and then the device just never got mailed out, just to get at, you know, this pure selection margin, which seems to play a large role.

So I think with a new technology, one has to be very careful about the difference between the ex post effects versus what the beliefs of consumers about this device would be, and I think that's where I would go on this monitoring piece.

So let me get to the part that I like more, just because I'm on the structural IO side of things, which is the equilibrium, the market for car insurance under different information regimes. And so there's been this whole advance in empirical IO especially coming from people analyzing the Affordable Care Act markets on trying to understand equilibrium in insurance markets. And I think a lot of what this paper does is take all that frontier and, like, put it into the auto insurance sector. And in some ways, the auto insurance sector is very compelling because in health insurance you have to deal with the fact that maybe I like Blue Cross Blue Shield because of the network or something like that, so there's all sorts of product differentiation.

For auto insurance, that product differentiation angle is just much less compelling. And so I think one can really kind of reduce things down to, like, the financial aspects of an auto insurance contract much more persuasively. And I think this is one of the -- when there's a talk about dimension reduction, I think this is what it's about, is that we can reduce -- we can reduce a whole bunch of driver characteristics into, like, an ex post utility with care preferences or whatnot.

Okay, so I think that's nice. It hits you with two problems. One is if all products are the same, then you have to understand why people are choosing choices that are completely dominated, that are just more expensive no matter what your accidents are. And then one of the pieces here is that you're going to have to account for people switching very infrequently. And I think this is not just like a little bug in the data that you have to kind of paper around. It's a real issue in the equilibrium in the market, right, which is as Sven Handel showed in his job market paper, if people don't switch that often, it kind of slows down the unraveling process in this -- in the equilibrium in this market, so it's not just fitting the data; it also changes the equilibrium. And I think this is a nice piece to put in there because it matters this way.

Okay, so some more comments. So there's a whole bunch of analysis in the paper trying to tell you that the model is doing a good job at fitting the data, and a large part of it is that there's some changes in, like, I forget the state changes its required insurance coverage from I think 30- to 50,000 or the other way around, and then you can say, well, in that state that we hold out of the analysis, what are the predicted versus realized market shares. And I think that's really neat.

It was hard for me to understand how much of that was coming from the change in the policy just not changing the averages too much, like the policy didn't change choices so much, or how much was coming from the model doing a good job at capturing the changes.

So it just -- I think it's a great idea to have this out-of-sample fit of the model. I just want to know a little bit more what I should take away from it.

And then just going deeper -- and this might be, you know, if one were to break up these two papers, there's this kind of idea of what should be information design in the auto insurance market. So right now, we have a very public record of all the accidents that occur, and you could imagine other types of organizations. You can imagine the firms keeping all that data private. You could, you know, imagine past claims kind of falling out after a couple of years from the information that firms could use. So there's a lot of policy design for this market that's relevant, even beyond this monitoring program.

And so I think there's a -- there's an amazing kind of policy discussion of how changing the disclosure of information on accidents changes the equilibrium in the market, making it public or making it private to firms, and I think that's very compelling. It's not something we thought about a lot. I'm always thinking that, you know, there's some countries that will stop kind of historical default information after, say, five years, and that changes the equilibrium in the credit market completely. And

I think there's a similar analysis here. So this is where I think it's a very powerful structure that you guys have put together.

Okay. And, yeah, so thank you for that.

(Applause.)

MR. KOCH: We have a couple minutes for questions, or if you wanted to respond.

If you have questions, speak out and we'll bring a microphone.

MR. JIN: So I actually prepared a very short deck. This -- out-of-sample fit is very well taken, this point. I will revise the paper, but given the time limit, I want to make sort of two clarifications and show you this analysis, which is Appendix G. I never thought it would, like, see the light of day, so thank you for that.

So the first clarification is that we focused on one-driver-one-vehicle polices, and that's actually quite important to making the analysis tractable, but I think there really is a lot to be done on those sort of multi-car-multi-agent sort of policies.

And two is that the finish rate, unconditionally, is 10 to 20 percent across state. Conditional on you starting, there is about 27 percent of people that drop out of the sort of -- just either never start or drop out, mostly within a grace period. So realistically, how does the firm manage the sort of, oh, maybe you don't know the program, maybe you don't know your own risk, is by just as you go along, I give you some feedback in terms of this is what I project your discount to be. I think that's also part of the reason why we start to see a lot of really sort of -- some of the advantageous action working so well is because there is this sort of mutual communication between the firm and the drivers.

So to your point, last point, about sort of thinking about the information content of different variables, here's the analysis that I did. So I simulated a risk pool that centered around, like, 4.4 percent accident risk, which is pretty representative, and this is the density of the true risk of each individual here. And then now on the Y axis I'm plotting instead of the density the firm's belief about this person's risk along the spectrum of true accident type, right? If the firm is amazing, is an oracle, then it will have this dotted, 45-degree line. Notice that I am, like, segmenting the Y axis and zooming in here.

But then what in reality the firm thinks everyone is the same risk because everyone is pooled together, right? So we have this flat prior here, centered at the mean risk. Now suppose I start to model, like, with some distributional assumption on prior, you can start to model how does this belief change over time as claim is being revealed.

So, of course, it's going to -- because claim is the sort of objective measure of risk, all right? Except that is very sparse, so as time go along, you sort of converge to the oracle. But this is what you -- this orange line is what you see with the sort of -- even just one period revelation of the sort of telematics or monitoring score. And you can see it's even more powerful for the safe drivers, which are really difficult to tease out because claims are so rare for them. So I think we can do a lot more analysis of this.

Another point is that, like, in the '90s, I actually saw quite a lot of papers about claim risk and disclosure because it's very difficult, even if people want to disclose claim, to enforce this data sharing. Like, how do I know you are sharing all of your claims with me, your competitor, right? So essentially what they end up doing is that they enforce -- so there is this industry organization called CLUE that goes into the back end of every single auto insurer. So as soon as you call to file a claim, this information will go to CLUE first before it hits the company. So I think with a lot of talk about sort of how do we do data sharing sort of more generally, I think this could be a useful precedent.

MR. ROSENBAUM: So I hope no one finds this deceptive, but in the interest of time, we're actually not going to take questions. You're more than welcome to chat with him after -- oh, one question, okay. I've been corrected. We have time for one question.

AUDIENCE MEMBER: Yeah, so, you know, one reason consumers might not opt in is if they prefer to keep their information private for reasons independent of selection on riskiness. They just value privacy. I wonder if there's any way to, you know, address, like, the impact of that and might there be a way to measure that, like say if there's some variation in whether the monitoring was time-limited or not?

MR. JIN: So you mean whether the data is kept for a limited amount of time?

AUDIENCE MEMBER: Well, like, suppose it was we're going to monitor you indefinitely versus only six months.

MR. JIN: Okay, yeah, that's definitely a big concern. So a lot of people ask why don't you do a counterfactual of continuous monitoring, and one of the things that we really can't say a lot about how -- sort of how much of that monitoring disutility term that we found on average $93, right, how much of that is really because of privacy concern because that's the part where -- or effort cost because you need to exert effort to be, like, monitored and to, like, appear safer, right?

So how much of that is really that cost? I think that's definitely focus of ongoing work, yeah. But that definitely speak to sort of there's a lot of -- a lot that goes into that term, the disutility term that we found.

(Applause.)

MR. KOCH: Just to note, the good news is the child care center is in the other building of the FTC, so in the event that we do turn PG-13 again, we should be okay and safe. That said, everything will be made on the record and put on the internet where there's no swearing, so please be mindful of that.

Our second paper for this session will be presented by Sharat Ganapati of Georgetown University. It's coauthored with Rebecca McKibbin of the University of Sydney, and it will be discussed by

Patricia Danzon of Wharton at completion of the talk. Thank you.

MR. GANAPATI: I'd like to thank the organizers and everyone here for selecting this paper. This is joint with Rebecca McKibbin, and it's a bit of -- it fits into my larger research agenda, which doesn't just look at a single country's context for monopoly but looks at how monopolies kind of interact and what we can learn from other countries in the context of both the U.S. and abroad.

So this is about the pharmaceutical industry, and, in fact, we're looking at a very specific point in the pharmaceutical industry, which are generic and off-patent pharmaceuticals. So this is motivated by this guy, Martin Shkreli, who's relatively famous for charging in the United States about $750 for a pill, which, you know, almost every other country around the world costs between $1 and $10 if you convert it from pounds or euros to the U.S. dollar.

And this is a drug where there's only one approved FDA supplier, and this man bought the rights or his company bought the rights to that drug and was able to charge a relatively high price in the U.S. Now, there are other drugs that are actually cheaper in America than in other countries around the world. Here’s another generic. It’s called gabapentin, and it’s used for epilepsy. It’s actually cheaper in America than most other countries. In the U.S., it costs about 17 cents a dose; while in most European countries, it’s more around a quarter a dose.

Now, if you look at it in the United States, we have over 20 approved FDA manufacturers for this drug. Well, in the U.K., you only have 11, and just -- this motivates kind of a big economic question, which is why doesn’t the law of one price hold. Now, as a trade economist, I think this holds, you know, a close part to my heart than most everyone else around here, but in this case, you know, there’s a few ways we can think about why the prices are not the same across the country.

The first is trade barriers. Now, if you look at pharmaceuticals, especially with First World countries, we have extremely low transport costs and tariffs do not bind, so that’s not a traditional explanation.

That brings us to kind of a bigger idea, which is the idea that fixed costs instead could play a role. Now, what are these fixed costs? Well, they could also be coming from an idea of imperfect competition, and that is going to relate to the idea of what generates these fixed costs. So you can get high fixed costs, and these can lead to very few entrants, which could lead to prices far away from kind of perfect competition. And these things can be driven by two things.

One is what I'm going to call entry barriers, so that's the FDA approval process; and the other item is something that is more fundamental to the market, which is some markets are just bigger and some markets are just smaller. So if you have a constant fixed cost, if you have a big market, well, you're going to get lots of entrants. If you have a small market and this constant fixed cost, you're going to get very few entrants and potentially higher price.

So this is going to read to kind of a bigger policy question, which we're not going to answer in entirety. We're going to just answer for a very small portion of the market, the generic pharmaceutical market, and that is why are only some drugs expensive in America. Not all drugs, but a very small subset of drugs are expensive in America.

So let's focus kind of from the big question onto what we're going to answer today, which is what is the role played by these fixed costs, and we're going to try to recover what is the cost of entering a market on market outcomes. And so this is going to matter for many contexts. It matters for trade; it matters for antitrust. If you have a very high fixed cost, there's not much that antitrust might be able to do and, in general, competitive law.

Now, in pharma, I know this is not a trade audience, but this is actually a big issue in future trade agreements that the U.S. is potentially negotiating or was negotiating as of two years ago.

And so this is also going to introduce a second set of questions, which is prices aren't just about market entry costs. And in a lot of contexts, especially in the pharmaceutical industry and in the medical industry, prices are not always purely competitive outcome; they're a product of some sort of bargaining or buyer/seller negotiations. So we're going to have to incorporate this type of pricing in a model where there are these differences in fixed cost.

And this relates to the larger question, is what happens to downstream monopsony. And so, you know, we don't always think about what this means in the medical situation, but in most European countries, we have a single buyer that is able to exert some sort of monopsony power and create certain market outcomes.

So I'm going to skip the literature here, and I'm going to get straight into kind of the data. So we're going to make a couple of assumptions here, and this is going to be applied more to the generic and off-patent market than it is to the on-patent market, and I just want to be aware of that, but we're going to look at these pharmaceuticals, which we're going to call nearly identical in every country. So off-patent, off-brand items are pretty much identical, but there are some questions of are medications in India and China, you know, not as safe as what's sold in the U.S. and the U.K., so we're just going to look at rich, English-speaking countries.

And so we're then also going to generalize away from the role of innovation because if you think about the pharmaceutical market, there is a role if, you know, we change prices, that's going to change the incentives to enter the market, we're going to generalize away from that. We're going to look at off-patent stuff, and we're not going to just look at off-patent pharmaceuticals; we're going to look at only those that are shelf-stable so you can have storage and also we're going to also not just look at things off-patent; we're going to add an extra five years' buffer after drugs go off-patent to kind of not worry about the initial market entry role, which is highly regulated in some markets.

We're not going to worry too much about what's called formulary design. We're going to assume that almost all of these drugs are available for consumers. We're not going to allow for kind of entry and exit of these. But even with this, even in this very, very simple kind of world, at least in my opinion a simple world, there are still many, many potential prices out there.

And so we're going to focus on a very, very specific subset of prices, and I'm going to first tell you what are the prices we're not going to use. We're not going to use what are available in these $100,000 data sets that are kind of wholesale prices before any lump sum rebates. We're going to also think about what happens with, you know, buyer copays and drug plan premiums, but at the end of the day, what really matters is the per-pill price net of all rebates, discounts, and dispensing fees paid by the combination of an end-user and/or the government or insurance company.

And, so, what we're going to end up doing is we're not going to look at private insurance in the

United States. We’re going to look at mostly public insurance markets where we have great price data, so we’re going to look in six markets. The United States, we’re going to look primarily at the Medicaid market. We’re going to look at Australia’s national PBS system. We’re going to look at Pharmac, which is the New Zealand system; BC Pharmacare and Ontario Drug’s benefits, which don’t cover the entirety of their populations but are kind of the public plans for two of the largest English-speaking provinces in Canada. And so all what these six markets are going to do is we’re going to kind of have a very specific set of prices that are going to be comparable across countries.

Now, for robustness, I’m not going to get too much into this. We’re also going to look at Medicare Part D in the United States and what we call the wholesale price, but I want to emphasize, we don’t actually observe the entirety of the price in kind of the context of comparison between countries in these markets.

So what we do with this data is we make it comparable across countries. That’s a quite large task, it turns out. Unit of observation is going to be a molecule dose form. And our key innovation here, actually, is we actually use public data. And I didn't realize this was a thing. I'm not a healthcare economist; my coauthor is. But almost every paper we saw does not use public data; they use a proprietary data set, which is really hard to kind of cross-validate and see what's going on because a lot of the pricing in these markets is opaque intentionally. And I do want to thank my RAs and coauthors for putting a lot of this together.

So just to give you a sense of what our data looks like, this is a comparison of in 2016 the drugs with the biggest price differences between the U.S. and British Columbia. And I'm just picking British Columbia because it's relatively close to the U.S. So if you look at this, we have a set of drugs here. They are all drugs that were invented relatively long ago.

So if you look on these, these are mostly drugs invented in the '50s, '60s, and '70s, though one example of a drug in the '90s is mebendazole, but that dosage is a new dosage. The actual drug's first approval date was with a capsule formulation, and that was in the '70s. So what I want to emphasize is everything here listed is old. This is not something that we need to worry about innovation.

And what we see is, you know, these are obviously the drugs of biggest price differences, so there are very few U.S.-approved manufacturers, and there are very -- relatively large price differences that we find. Now, just to kind of show you what all data we have, again, comparability isn't perfect, so we have different ranges of data for different markets, but in general, the U.S. is a higher price than foreign markets, and we're looking at markets that have a variance in the number of potential manufacturers in the U.S. But on average, we have about four manufacturers entering the U.S. market.

Now, one key fact, and this key fact drives our entire analysis, is we can look at the number of U.S.-approved suppliers, which is on the X axis, and we can look at the difference between the U.S. price and the foreign market price as a function of how many firms got U.S. approval to enter the marketplace. So if we look at just drugs with just one supplier in the United States and compare it to Australia, British Columbia, New Zealand, or the United Kingdom, we have about, you know, 300 log points increase in the price in the U.S. marketplace.

And that is a log linearly -- semi-log-decreasing function. As you get more and more entrants in the United States, the price differential from the U.S. markets converges quite rapidly to foreign markets. And by the time you get seven or plus manufacturers, which I've used as the omitted set here to normalize the data, you're effectively at the same price.

And, so, this is looking at Medicaid data. This holds for Medicare data. It holds for MDAC data. It doesn't really matter what data you look at. You get some sort of downward relationship that is super robust.

And, so, another thing that's going on in this medical marketplace, and in the interest time, I'm not going to go through the full kind of details, is we also find that generic drug demand is inelastic. And this is because of one thing we feel is, you know, maybe not everyone shoulders the full cost. And this is, you know, very common in Medicare and a lot of a foreign systems, but we can also try to actually show this in this one context because one nice thing about the wholesale drug marketplace is most of these drugs are not actually made in the United States. And so if they're not made in the United States, they're often made in a foreign country, and we actually have data on what country these drugs are made in.

And so one thing we do is we can actually say, hey, we actually have a cost shifter. And this cost shifter varies on the different drugs because some of these drugs are made in China, some of these drugs are made in the Philippines, some of these drugs are made in India. So we have these exchange rates. Our simplifying assumption is that we're going to assume that exchange rates are not functions of medical demand, and I think that's a relatively straightforward assumption to make. Exchange rates are changing for other reasons, and we can show that, you know, prices -- changes in prices don't affect how much we're paying for -- or how much we buy these drugs.

So with that idea, we're going to figure out kind of how to do a pricing model. We're going to have this inelastic demand, but we also have some key facts that we want to explain. And, so, we're going to have a few key elements we want in the model. We're going to include the roles of kind of suppliers, competition with the suppliers, but also the role of kind of like the downstream buyer.

In the background, and I'm not going to talk too much about this today, there's also going to be a competition between a branded drug and the generic entrant. And I'm not going to talk too much about it. It's in kind of the underlying part of the model, but there is a kind of second competition we have to also worry about.

Now, we have desires of this model. We want it to be simple, and I'm an IO economist, I'm as guilty as everyone else here in making a very complicated model with more bells and whistles than was really necessary, but we also do want it to be flexible. And a lot of these papers we've seen in this kind of drug marketplace are hyperspecific. You've got millions of fixed effects to kind of estimate, and it's not clear whether you're just picking up noise and what the validity of those things are. So we're going to try to take a model, we're going to try to use as much IO as we can, but we don't want to be too IO-y in some sense.

And so what we're going to do is we're going to kind of think of this first in terms of what price we're doing, and then I'm going to show you kind of what the model is we're doing. So when we look at prices in this marketplace, prices are actually very, very complicated. A price is a function of a pharmacy's markup, what we call a pharmaceutical benefits manager. That's some sort of middleman that does a lot of price negotiation. There's a wholesaler in the background. There's the manufacturer's markup, and then you finally get to kind of some sort of underlying marginal cost. And, again, even this is a simplification of the overall marketplace. You can find other players that have their own cuts of all sorts of the marketplace.

Now, we're just going to kind of compress all of these markups into a single markup over the entire value chain, and we're going to consider what that role of that markup is. And so in some sense, this is all that really matters for welfare if you don't worry about any sort of externalities that are imposed on the marketplace by all these intermediate players.

So this is, again, a simplifying assumption, but this is also kind of the problem with what data we have. If you don't have data at any intermediate stage, it's unclear what we're picking up at markups at different points. So we're going to compress all of these markups into one.

So we're going to have a two-period game, and this game is going to be relatively straightforward. There's going to be an entry stage, and there's going to be a price competition stage.

The entry stage is generic suppliers are going to choose to enter the marketplace. They're going to pay some sort of fixed cost. This fixed cost is going to have lots and lots of potential components, and we're not going to be able to disentangle all of those components. They can be rearing from everywhere from political interference to regulatory cost to bilateral payoffs to downstream prescribers, for example, to doctors.

And one thing I want to emphasize here is we're going to essentially bound kind of what these fixed costs are, which are the profit or the marginal operating profit of the Nth or Fth supplier in the marketplace. And another thing we're going to assume is market entry costs are going to be independent through countries. And that seems a little weird, right? I mean, in the on-patent marketplace, we would never make that assumption because there is a fixed cost of developing these drugs to testing.

But in the generic marketplace, it's actually very different. So one thing I did is I actually had an RA go through and try to count at least for a sample of the drugs the number of potential factories that have FDA approval or an equivalent approval of a similar First World country and a Third World country that can make these drugs. So this one drug I present at the very beginning, Daraprim, has only one approved manufacturer in the U.S. If you actually go to India and you look at the number of factories that supply it, there are at least ten, and in China, there are at least 62, and the vast majority of these factories make at least one other drug that has FDA approval. So it's not hard for them to get FDA approval if they wanted to for those drugs.

So following this market entry stage from this mass of unlimited potential suppliers, we're going to find there's going to be a subset that are going to pay a fixed cost. After that subset pays a fixed cost, they're going to choose a price. And that price is going to be negotiated with a final buyer, and we're going to assume kind of a monopolist final buyer in this case.

So there's going to be a kind of a profit that's going to kind of be function of a markup that's negotiated, some sort of marginal cost, and a quantity of drugs they sell. And I'm going to be specific here. We're going to be agnostic on the type of actual competition there, but we're going to recover the type of -- the resulting prices that are a function of kind of market characteristics because we have lots of different observations in this marketplace. Following this, sales are made. And, again, we’re going to focus on public plans with mostly inelastic demands.

So how are prices negotiated? So I’m going to focus on first one case of a monopolist buyer and monopolist seller, and then we’re going to be agnostic after that. So what are we negotiating here? Well, we have a Nash surplus, which is going to be a function of some sort of surplus of the seller, some sort of surplus of the buyer. So the seller is trying to maximize price minus cost. I’m assuming their outside option is if they don’t sell it, they get zero. And for the buyer, this is a little harder. If there is a monopolist buyer and monopolist seller, if negotiations break down, what happens? We need to figure out what’s going on.

So what I’m going to do is we’re going to introduce this concept and we’re just going to call it a choke price. This is some sort of a negotiation price if negotiations break down. So if there’s only seller of a drug and the buyer can’t come to an agreement with the seller for a price, there is some sort of residual price that’s going to be charged. This could be because there’s a political pressure.

There's also compounding pharmacies. There are a lot of kind of outside options. We don't know what those outside options are, and we're going to actually recover what this choke price is.

So the first order conditions in this kind of Nash setup are pretty straightforward. This is kind of from your intro to any IO type class. You get a monopolist price that's going to be a weighted function, depending on the bargaining weights of two things -- the marginal cost and the outside option of the buyer.

And that's a pretty straightforward kind of thing, which has two corner solutions. One is if you have perfect competition, you get price equals marginal cost. If you have a kind of all the bargaining weight on the seller, you have a seller with kind of perfect ability to extract out all the surplus. The price equals whatever the choke price and they extract out all the surplus from the buyer side. So you get a range of two prices here.

Now, what happens if there's more than one upstream seller? So I gave you kind of the baseline scenario where you have one seller and one buyer. But there are cases where you have multiple sellers, as I point out in the data. Well, what we're going to do is we're actually not going to take as close a stance. It's going to end up looking very Cournot-like, but it's not exactly Cournot, which is there's a function that literally just maps the number of -- the set of sellers to a set of markups. So what we're going to say is if you have seven sellers, for example, we're going to empirically recover that the markups are 30 percent or something along those lines.

And so what we're going to do is we're going to weight between the Nash solution and kind of perfect competition in this not -- well, nonlinear way which we're going to actually end up putting some sort of form on, but we're going to weight kind of you can have this monopoly outcome or you can have a perfect competition outcome, and where you are between those two outcomes is entirely dependent on the number or the intensity of competition.

So I want to emphasize we can take the setup and I can give you a functional form that is the same as either Bertrand or it's the same as Cournot. There's many, many variations of it, but the entire intuition I want to raise here is conditional on the number of entrants, pricing is fully determined in the marketplace.

And for tractability, at least for the talk today, we're going to do some things here. We're going to assume that the choke price is some sort of multiplicative function of the marginal cost. That is an assumption. We can try to think about how we can generalize that assumption, and we can also parameterize competition. This is effectively taking almost a Cournot stance, which is going to be a function of alpha, which is a parameter we're going to recover, times the log number of competitors in the marketplace.

So let's kind of take a step back from the pricing set into looking at what's happening on the market entry side of the problem, and that is we're going to look at this concept of excess profits. Now, what are excess profits? That is going to be if we have constant marginal cost, which is an assumption we're going to make, how much more operating market is it to take -- will it take to enter one market versus another marketplace. So how much more excess entry and profit will it take to enter, for example, the EU market versus the U.S. market and so forth.

And so, you know, we have data on only three market sizes. We're going to look at the U.S., the U.K., and Australia, and we're going to kind of make a comparison of how much more does it cost to enter the

U.S. market for a particular drug versus other markets that we see in our data, this -- emphasis on Australia and the U.K.

And what this is is literally a pretty straightforward thing. We take up the marketplace. We divide up the -- kind of the operating profits between all the entrants, and we see how much more it costs to enter the U.S. than a foreign marketplace. And I want to emphasize this is only done for the marginal generic entrant. We're not doing this for kind of Pfizer has a drug that goes off-patent, and so, like, so they take Viagra, that goes off-patent, we're not going to look at kind of Pfizer's incentives; we're going to look at the marginal generic companies' entrance rather.

And we can do very straightforward bounding exercises with this, how many more entrants could the U.S. support if the U.S. fixed costs were in line with other countries around the world, and we can take that and take kind of a welfare analysis of that.

So just to go -- I'm not going to go through the full estimation here. I'm just going to tell you the results and focus on the first column, which is looking at the Medicaid market in the United States. We find competition binding, but we also find that what we get is we get bargaining in many markets from Australia to the United Kingdom which look very, very close to a perfect buyer that effectively goes to take it or sell it off.

So what this cap -- or this first term is, this bargaining term, if it equals one, they're perfect -- perfect bargainers. They can extract out all the surplus as in terms of the buyers. If this term goes close and closer to infinity, that puts all the bargaining weight on the seller of the drug. So in the United States, we have sellers that have relatively high bargaining weights. And, again, this isn't a weight; this is a transform of the weight from 0 to 1 to 1 to infinity, and that's just a way of getting at the data.

We find that the U.S. just looked pretty terrible in this sense. And then we can take this data, feed it into kind of a market entry stage. We can look at how many million dollars in a flow million dollars per year does it cost to enter the U.S. And it turns out if you're comparing the U.S. to the Australian market or U.S. to the U.K. market, we get a cost between $5 to $10 million a year for the average generic drug.

And that seems low or high depending on your priors, but let's take this and kind of project it onto overall spending, at least with public plans in the United States to see what happens. And we're going to do a few counterfactuals. So the first counterfactual we're going to do is there's lots of variation in the number of sellers, and we're going to do a very simple idea, which is if it's profitable in one country, that drug or that maker is allowed to sell in every other English-speaking country because the labels are supposedly the same.

And so we're not going to change the market entry incentives. We're just going to say -- we're going to exogenously increase the number of sellers. So, for example, if there are eight sellers in the U.K., three sellers in the U.S., well, those eight sellers can also sell in the U.S. at no extra fixed cost. But we're not going to change entry and exit.

And so with that, what we get is we're going to look at the cost savings in Medicaid, and we find about an 8 percent cost savings on generics and off-patent drugs in Medicaid if you do that policy.

We can do a few other policies. One is that what if bargaining in the United States looked like other countries, so looks like the United Kingdom? We get a cost savings of about 20 percent. Now, we can combine kind of the single-market effect and bargaining. Well, it turns out it doesn't matter because once you start bargaining like other countries, well, you're already giving a take-it-or-leave-it offer, so you're extracting out all the surplus, there is no difference.

But, lastly, we can do finally something which is what if we changed the free market entry condition in the United States to look like every other country. And there is going to be some integer constraints here, but in general, what you get is you get a very similar cost reduction. In our empirical case for 2017, you get a 16 percent cost reduction, which is almost identical to the bargaining outcome.

So what we do is we take this to kind of imply that, you know, integer constraints do bind in some sense, but in general, you have two kind of ways of reducing at least drug prices in these markets. One is a very free market approach, which is reduce entry cost, and that is something that has, you know, been talked about by the FDA and a lot of regulators in a lot of countries, how do we make it cheaper to enter our marketplace.

The other option is take a United Kingdom/Australia approach, which is you only let one or two sellers in but you give them really binding take-it-or-leave-it offers on the table. And as long as you have an epsilon over kind of marginal cost, the sellers will take those take-it-or-leave-it offers and you can increase a kind of -- or decrease overall spending on pharmaceuticals.

So with that, I just kind of wanted to show that, you know, this is a project that, you know, takes a very complicated drug market and tries to simplify it down to try to distill out two core things that can go on. And those two core things are kind of policy-relevant, which is do we negotiate drug prices; and the second policy thing is do we allow free entry to show at least in one context they're actually relatively equivalent policies and become -- it kind of falls on the policymaker to kind of decide which is more politically feasible and implementable to go on from there.

Thanks.

(Applause.)

MS. DANZON: Okay, thank you very much for inviting me and thank you for a very interesting and provocative paper. It's an ambitious paper. You've just heard all that went into it. A brief overview is that what's being done here is to estimate the price ratios focusing on generic and off-patent brand drugs for -- estimate the price ratios for the U.S. market relative to five other English-speaking countries, all of them notably smaller markets. And the conclusion here is that U.S. prices are significantly higher if there are fewer than six sellers.

I would point out that the conclusion was overall the U.S. prices were significantly lower if I read Table 1 correctly. But focusing on the products that have fewer suppliers, the U.S. was more expensive. And then these empirical estimates are used to estimate parameters of the structural bargaining model between a U.S. payer and a generic supplier.

And then that structural model is used to estimate the effects of two policy changes in the U.S. -- reciprocity of approvals, which is equated to removing nontariff barriers, so reducing fixed costs of entry, and then federal bargaining over prices. And the conclusion is federal bargaining over prices would be more effective than -- because adding -- removing nontariff barriers over and above that adds very little.

So overview of my comments. I am seriously concerned that the price measure used is the reimbursement price paid to pharmacies by Medicaid and for a couple of reasons that I'll explain that this overestimates the actual price received by generic sellers. And since this is about -- the paper is really about the effect of competition in the seller market, I do think that if we're not observing the seller price that is potentially important.

If we're talking about overall policy, the fact that the sample of drugs is certainly not representative of the overall market is important. It's focusing on those products that are really quite old, and so in those markets having few sellers may be markets where, in fact, there's been exit, and so they're not typical.

The structural bargaining model, I think, does leave out some very important portfolio effects I'll elaborate on. I'm not so sure about the lessons from foreign markets, and so I'll talk about what policy implications I think we can look at here.

So, first, how are generic prices determined in the U.S.? As Sharat explains in the paper, basically the pharmacists can substitute between AB-rated generics. That means the generics that have the identical molecule dosage form and strength and have been shown to be a bioequivalent, and so the decision-makers, the buyers for pharmacists -- for pharmaceuticals are the pharmacies.

The private payers represented usually by their PBMs, their PDPs, they reimburse the pharmacies for generics based on a MAC, a maximum allowable cost, and the point of that is that that pays a uniform amount for all equivalent products, all substitutable products. And that creates an incentive for the generic suppliers to compete below the MAC because the pharmacy keeps the margin below the acquisition cost and the MAC. That becomes a confidential rebate or profit to the pharmacy, and then periodically the payers audit the pharmacy acquisition prices and reduce the MACs to recoup the savings from competition but with a lag.

And so the private payer price to the pharmacy overstates the generic supplier price by the amount of the rebates that are being given to the pharmacies, which are nonobservable.

Now, the price that's actually being used in the paper is not the private payer price but the Medicaid price, and Medicaid is about 10 percent of the market. And under the Affordable Care Act, the Medicaid upper limit price, which is generally what is used, is 175 percent of the average weighted average manufacturer price. The average manufacturer price, or AMP, is the price we would ideally like to measure because it is the price received by the sellers, net of all rebates given to pharmacies. But that is unobservable, and so what the paper uses is the Medicaid reimbursement price, which is 175 percent of the AMP.

States can choose to use a lower MAC for Medicaid, but that's not the norm. They argue that -- pharmacy associations argue that that would put the independent pharmacies out of business, which would not be good for Medicaid beneficiaries. And so what's being used is Medicaid reimbursement, which represents 10 percent of sales in the U.S. And it's based on this FUL which exceeds the private payer price, and that exceeds what is received by the sellers because of the generic rebates that go to the pharmacies.

So that's one concern. Second concern is including only the oldest products in the market. So only the generic markets that are at least 20 years from the FDA approval of the originator product are included, but that includes generics that have come to market relatively recently. And, indeed, the median, I think, or mean date of FDA approval of the products in this sample is the early ‘80s, so we’re looking at really old drugs.

And typically in a generic market, you’d start off with a few suppliers and the number would increase, and then there will be exit. And so my concern is the markets we’re looking at here with few suppliers in many cases would be markets where exit had occurred because the market had become unprofitable.

So in that case, you know, I think we really need to understand what it is that is bringing about small number suppliers. Is it just relatively small markets? Is it relatively high fixed costs because of the age of the market, because it is true, technology changes rapidly in this -- in the manufacturing of generic drugs. So if you brought your product to market 20 years ago, that is very out of date for current manufacturing techniques, and so there could well be big retrofit costs of staying in the market. So, you know, what the costs are for those particular products, I think, could be quite different from an average.

How bargaining actually works in this market, I think it’s really important to understand that it’s the pharmacies that are the purchasers here, not the buyers. The pharmacies in the U.S. market, as we all know, are huge chains. They are bargaining with the generic suppliers. They're bargaining from a central corporate headquarters for the entire portfolio of products for all the chains, all the stores in their chain. So think of it as headquarters of CVS Caremark bargaining with the generic suppliers, so they set it over the entire portfolio.

And so what they're looking at is obviously lower prices, but it's also the breadth of the portfolio, it's how many of the newest products that are going to come to market with that big margin on the 180-day exclusivity -- I won't go into the details of it -- but those are some of the new products come to market with a big potential margin. That's very important to the pharmacies.

And, also, the big generic suppliers provide restocking services. They monitor when individual stores need restocking, and reliability is also important. So the notion that there's just a fixed cost to pay and then an entrant could come in and actually supply this market leaves out all the other factors that the customers are actually looking at, which is breadth of portfolio, reliability, and all of those factors. So leaving that out I think is potentially important in thinking about what the benefits of entry may be.

I think that there's a mischaracterization of this sort of magical bargaining power that the foreign payers are using because actually most of them are using something very similar to what the U.S. does. The Canadian provinces, it is true, use a percentage of the originator price, where that percentage depends on the number of generics in the market, but as a result of this, there's a lot of concern in Canada that the payers are not actually capturing the discounts that are being given by the suppliers to the pharmacies in Canada as they are in the U.S., so that the payer is not recouping the savings from price competition as the U.S. payers do because of the MAC being adjusted.

In the U.K., in Australia, what they're actually looking at is market prices and using a sort of similar system that's very similar to the MAC used here. Australia calls it reference pricing. The MAC is a form of reference pricing. New Zealand does do competitive tenders, but only for particular therapeutic classes. New Zealand is a tiny market. I think last time I looked the population of New Zealand was a bit bigger than Philadelphia, so, you know, you can supply the New Zealand with one or two suppliers. You cannot supply the U.S. reliably with one or two suppliers, so it's a very different situation.

So policy options, I'm concerned that in the modeling of the need for and the effects of federal bargaining, the federal government would not be able to walk away from particular suppliers the way New Zealand does because U.S. consumers count on reliability and availability of all the generics. So I really am not confident that tendering by is feasible and I think the bargaining that's being done by the big pharmacy chains is probably as effective as what's being done in other countries.

Reducing the tariff barriers could indeed certainly reduce regulatory costs, but I wonder how much of the actual barriers are related to these portfolio issues, which wouldn't be affected by regulatory reduction.

Finally, I think alternatives that would be worth looking at are federal limits on unreasonable price increases when there is either a changeover of ownership or exit. That is, in fact, when we see these big price hikes. And, so, you know, a more surgical sort of policy that would address those issues, I think, could be considered.

And, finally, obviously, if we get down to one or two suppliers, there is a big role for antitrust enforcement as, in fact, is actually happening in the market right now. But for the market in general, I think the U.S. market is functioning remarkably well. Thank you.

(Applause.)

MR. KOCH: So we do have time for questions from the audience, so if you do have a question for Professor Ganapati, please reach out to someone with a microphone.

MR. GANAPATI: Well, I guess I'll just say a few things if there's no questions. One thing is I want to emphasize we don't think that bargaining is ever happening in the marketplace in the explicit sense. This is a modeling construct. We used the bargaining to kind of simplify what's happening in the world, and I want to emphasize when, you know, we say there's bargaining in the world and we take out these Nash surpluses and think of what's going in the marketplace, there are many, many different ways to generate data. This is just kind of an assumption on kind of how the data-generating process works.

And I do want to emphasize, you know, the discussant is entirely right. We're not looking at every drug in the marketplace. We're looking at a very subset of selected drugs. And so we're not trying to say that, you know, this solves all of America's drug problems in, you know, one sentence. We were looking at -- and these older drugs, there are some -- for some reason, you know, 20 sellers in Europe for some of these markets and only one in the United States and trying to figure out why are there this. Those fixed costs represent kind of the cost of setting up a marketplace in the United States and includes setting up kind of reliable transportation, reliable kind of bargaining with the CVSs and bilateral payments of all sorts.

And we're not taking a stance on what goes into that fixed cost. It's a large fixed cost, and all of the things you've mentioned are part of that fixed cost, and it's kind of a future work to kind of figure out how to disentangle what's going on because you have everything from the U.S. being a bigger country with more to distribute to the fact that there's different demands for different drugs in the United States.

AUDIENCE MEMBER: Hi. I guess one way to -- there's two ways to think about your results, and one is that there’s something different about the distribution of fixed costs in the United States from other countries. And the other is to say there’s something different about the elasticity of demand for drugs in the U.S. versus other countries. And it seems like you’re leaning towards the fixed costs explanation, but, like, do you have a sense of what’s in that and why that is? Like, usually we think of, like, opening a business and things like that, and regulatory approval are high in the U.S., but we often don’t think they’re lower in European countries.

MR. GANAPATI: Yeah, so my coauthor talked to a few regulators, both in the U.S. and abroad, and in most countries, we agree that in most industries the U.S. should -- seemed to have a lower fixed cost, but that does not seem to be true, especially in the pharmaceutical industry, and that is a mixture of everything from higher costs to just set up the distribution networks, to negotiating with a small -- negotiating with, you know, three buyers is still harder than negotiating with one buyer.

And there are so many kind of back-channel bilateral payments, so if you talk to a lot of these drug makers, they’re saying essentially we’ve got to pay off CVS to carry our product. And that creates all sort of externalities. We don't know what those costs are. We're just kind of lumping them all into our fixed cost.

MS. DANZON: One question I had, Sharat, I don't know if you looked at it, was you talked about standardizing the products so they're the same molecule form, but form could be just tablet, right?

MR. GANAPATI: Yes.

MS. DANZON: And so certainly when we looked at the data, we found a lot more of the relatively expensive delayed-release forms of tablets, which you do not find in other countries because they get some sort of IP in the U.S. where they don't in other countries, and so if you've got that sort of blending of forms within your tablets, that could be one of the factors that's happening, especially for those smaller drugs.

MR. GANAPATI: Yeah, so it's not in the current draft of the paper, but we have another analysis subset, two drugs without extended-release forms. So just tablets without extended release, we find almost identical results.

MR. KOCH: We will now break, I believe half an hour until 11:00, when we will have a keynote address. Thank you.

(Applause.)

(Recess.)

KEYNOTE ADDRESS: SCREENING FOR PATENT QUALITY

MR. KOCH: Good morning. We will now be moving on to our next speaker. Our next speaker is Mark Schankerman. He's a professor of economics at the London School of Economics and a research fellow at the Center for Economic Policy Research in London.

He received his Ph.D. in economics from Harvard University, formerly taught at New York University, and was a research associate at the National Bureau of Economic Research for ten years. He has contributed extensively to the literature on patents, research and development, productivity, and the economics of emerging economies.

We look forward to your talk. Thank you.

(Applause.)

AUDIENCE MEMBER: (Off microphone comment.)

DR. SCHANKERMAN: No, but, you know, as we know, building on the -- standing on the shoulders of giants who suffered more.

First of all, thank you, everybody, for being here, and I know that it's late in the two days. I want to thank the organizers again for inviting me to be on the committees, on a different committee, and participating in the conference. It's my first here at the FTC, and I'm hoping optimistically that it won't be the last.

Today, what I want to talk about is something completely different, to quote Monty Python, which is patents, screening for patent quality. Now, this work, which, by the way, is under revision for a journal and we've been revising it for a year and a half, our deadline is next May, so hopefully we will be done by then. This is joint with Florian Schuett, who is at Tilburg University in Holland.

So in 1999, Amazon got a patent on one-click shopping, as you know. And you probably all know that this was a highly -- well, this was a patent which allowed you to complete as a customer a transaction without having to repeatedly enter your data, your customer data. And by all accounts, it was highly profitable. Nobody's been able to measure the profitability, but by all accounts it was highly profitable, and that's why it's famous.

At the same time, when it was issued, many observers, perhaps even most, many observers commented that they were extremely skeptical that this thing should ever have been granted. Not that it wasn't valuable, they all recognized that. Not that it wasn't necessarily creative, it might have been, good idea. But that doesn't pass patentability standards as I'll talk in a moment -- talk about in a moment.

And yet even though many skeptics thought that it would not have passed so-called nonobvious -- novelty and nonobviousness requirement for patents -- you can't do something that's too close to something else or that would be obvious based on what else has been done prior. It was never challenged in court. And in 2017, it expired after full term.

Okay, so here's a patent, highly valuable, questionable in validity in a sense of patentability requirements, but it never got challenged, okay? This patent actually illustrates some of the core things I want to talk about in this -- in this talk and what we're trying to do in this paper.

The central issue here is is that typical, or is that an outlier? Well, more generally, how bad is the so-called patent quality problem? There's a lot of discussion in the literature, particularly in the law and economics, legal scholars, discusses all the time, Congress has stepped in with the American Invents Act in 2011, which was the most important probably for 50 years, most important piece of legislation in relation to patents.

The Supreme Court has stepped in on a number -- in a number of very high-profile cases, notably

Merck about ten years ago, but others as well. And all of these had the consequence -- they were all worried about the patent quality problem, that this proliferation of patents, many of which, most of which, who knows, shouldn't have been granted, so the conventional wisdom goes, and they've narrowed patent protection or the enforcement or enforceability of patent protection. And I can talk about how if you're interested, but the consequence has been to narrow or weaken patent rights. So there's a big pushback against patent rights.

Now, the question is how bad is the patent quality problem, okay? As the well known economist in the 19the Century, Mark Twain, once said, you know, it ain't all the things that we know -- that we don't know that are dangerous; it's all the things that we do know that ain't true. Okay, so is this conventional true or isn't it true? And how serious is it? That's really the question here.

Now, before I get started with this, let me just say, there are, I think, implications for the Federal Trade Commission and competition policy. Maybe we can take them up in discussion, but the most important and obvious one is that the market for technology, that is the diffusion of technology, i.e., licensing or sale of patents. This market is going to be undermined by asymmetric information, the standard bargaining problems that can arise, but one of the pieces of information that may be -- one of the things that may be very uncertain is whether the patent which you’re asking me to pay a royalty on is likely to be upheld if I challenged it. So having a patent quality problem creates a licensing problem, and that may create licensing-connected competition problems, not least of which, of course, is the alleged trolling behavior, which we’ll come back to in a moment. So I think there are links to the interests of perhaps more people here.

Now, what should we do about all this? Well, some legal scholars -- Lemley in particular most famously at Stanford -- said, look, here’s this rational ignorance argument that says don’t worry about it, okay? Don’t worry -- what we should do is basically let the court sort this out. And the argument is that most patents are not valuable, that’s true. My own work on patent renewals and others from all that stuff we know very well that that’s true.

It’s also true, as he says, that a very small fraction are ever litigated. He says 1 percent in that paper; it’s more like 2 now. And he said, therefore, why spend a lot of money on Patent Office examination if most of this stuff never gets challenged anyway? Okay, now, my view is that this is fundamentally wrong. I have to say in Mark's defense -- Mark Lemley's defense, when I challenged him, he said, well, I didn't really mean that, so okay, fine.

But in any case -- in any case, I think it's fundamentally wrong, and the reason is twofold, and the Amazon case illustrates it. First, it's wrong because bad patents -- by bad I mean those that shouldn't have been granted, and I'll be more explicit about that in a moment -- that bad patents may not get challenged, and, in fact, we'll see the probability that they are challenged from the model, from the simulations we do. That's number one.

Number two is if they aren't challenged, they shouldn't be charging royalties or getting royalties, but they are, and, therefore, prices are higher. Okay? So the fact that they don't get -- that not many are challenged is not a reason to just not worry about examination. In fact, I'll argue in this talk that that's exactly why you should focus on examination as opposed to the courts.

Okay, now, so the questions are the first bullet, which somehow got erased, is how severe is the patent quality problem. The second, should examination be intensified? It's expensive to do that. Should we intensify it, or should we go the other way and just have a registration system like copyrights? There's no examination of copyrights. They last a very long time. Maybe we should do that with patents. Or should we just move it all to the courts like the rational ignorance argument of Lemley says?

Second, they charge -- the Patent Office charges lots of fees. They're not huge, but here they are. The current Patent Office -- U.S. Patent Office -- to apply for a patent, there's a whole set of fees. This is summarizing them, is something on the order of $2,000. It could be a little more depending on the number of claims. If you -- then you have to pay after you get a patent granted. You have to pay renewal fees to keep it in force. If you don't, it expires, it lapses, up to 20 years. And if you pay all of them undiscounted, it's about $14,000. Okay? So there's a nontrivial amount of money. This is per patent.

Now, currently, most of the fees are post -- are post-grant, the renewal fees. The application fees are low. Is that structure right? Should we change it? Should we backload all fees? Should we frontload all fees? What would that do to screening? All of these questions I'll try to answer, okay, both theoretically and empirically.

And there are other things, too, that have happened. For example, the Patent Office after the 2011 America Invests Act introduced, or strengthened really, but essentially introduced a post-grant review. And what that is is that anybody can challenge a patent application after it's granted but within the Patent Office system. So it's like a prelitigation litigation. It's much cheaper. It's probably less accurate than the courts arguably. Is that a good idea? Many people complaining about it now, but is that a good idea or not? Does that improve screening or not?

So what I want to emphasize here is what's motivating this whole project are these policy questions. Okay, now, these policy questions can't be answered by empirical work, and they can't be answered by a structural model with empirical work because there's no variation in the -- well, it's obvious, okay -- if you're looking at one country, but what we do here is to build a structural model, if you want to call it that, or a model, and then we're going to fake empirical work, i.e., we're going to do calibrated simulations, okay? But they're calibrated on real data, so it's a poor man's version of empirical work.

Okay, so what we're going to do is we develop a model, a patent screening, and one of the important things about this I really want to emphasize, it's more important than the technical things, is the patent screening isn't in the Patent Office only. This is the fundamental point, that because there are lots of comments by legal scholars, suggestions of how to improve this or how to improve that, but they're not looked at as a whole, and they're not looked at in equilibrium.

And when I say that there's a whole system of screening, I mean there's first the self-screening into whether you apply for a patent. Then there's the screening within the Patent -- the administrative screening within the Patent Office itself. Then there's -- and then, after that, there's the screening potentially by the courts if it gets to a court. Okay? But the intermediating -- the stage between those two -- the patent grant and the court -- is the licensing game, because maybe I can license in such a way that keeps you from challenging me, i.e., Amazon, because that's exactly what happened between Amazon and Barnes & Noble.

So we want to have this more holistic view of screening, and we want to embed it in an equilibrium framework so that when you -- we can look at the instruments and see whether there are unintended consequences of playing with these instruments, these policy instruments. So that's the objective. That's the objective here.

Now, the way we're going to do this, we have to build a model, and the model's going to be simplified obviously, but we hope realistic -- reasonably realistic. So in this model, there's an inventor, and this inventor has an idea. The ideas are exogenous, so we don't model the supply of ideas because I don't know anybody -- I've been working in this field for years, and if I don't know how to do that, I don't think anybody does. But we don't want it to be contingent on that, so that's given.

The inventor has private information about whether his patent's valid, that is, should be granted, and I'll give you the criteria in a moment. The competitor doesn't know this. The single competitor doesn't know this, but he updates beliefs about the inventor's type, valid or not valid. I'll call it low and high type, okay? And he updates when he sees the various actions of the Patent Office, whether it's granted, whether the -- if it is granted, whether you pay the fees to keep it in force, and so on, and also sees the license agreement that you offer to him after you get a grant. So all of this contains some kind of information he based in updates.

Now, the Patent Office and the courts receive an informative signal about validity, if you want, okay? The Patent Office -- the key thing to realize is the Patent Office, by law, screens everybody. There's no selection once you've applied. Everybody gets screened. And -- but we're going to model that as an imperfect signal. So the Patent Office is going to make mistakes. So they're sometimes going to grant invalid patents, but they're always going to grant valid ones. We can have two-sided errors, that doesn't change anything here. Almost all commentators think that the problem with the Patent Office and screening is that they don't grant -- they grant stuff they shouldn't rather than they don't grant stuff they should. So that's how we're modeling it in the baseline.

The courts, on the other hand, get a perfect signal, that is, they don't make any mistakes. Now, the reason -- it's not that we believe that, but we want to give as much to the rational ignorance argument as possible. We want to say let's let the courts make no mistakes, okay? Can we still -- how much can we rely on courts as opposed to the PTO? Okay? So, again, we can -- we've generalized all of these things in the paper.

But the key difference here, the courts have the advantage of making no mistakes, but they only judge those cases that get to them. So they never judged Amazon, okay? And that's the difference between the Patent Office and the courts.

Now, in this framework, we're going to be able to look at all the instruments in question that are available, and the instruments are going to be the Patent Office fees, pre-grant, post-grant, the intensity of examination within the Patent Office, and some other things we'll talk about and look at those in a framework in which all of these things -- all of the outcomes are linked because there are going to be various interactions.

And then we're going to parameterize this model based on actual data, and I'll try to get to that. I hope I have time.

So let me just give you a quick summary of the results. First -- no, sorry, I have to advance it, yeah.

First, frontloading fees improves welfare and improves screening. So we should be moving all the fees from post-grant to pre-grant. I'm aware there's an income constraint argument against this, but I'll come back to that if you want, you know, that maybe some small patentees can't afford if we do that, but we can come back to that. But that's a big -- that's going to make a big difference we'll see in the simulations.

Secondly -- okay, secondly, the courts, even if they're perfect, as we're assuming in the baseline model, they cannot achieve full screening. Full screening means they screen out all the bad types. And the reason is many of the bad types never get there, okay, the Amazon case again.

And, finally, and I'll just say that a little bit later hopefully in a little more detail, the incentives to challenge, because you think it's an invalid type -- invalid patent, are inefficient. And it's not just for the reasons that others like Carl Shapiro, Joe Farrel, and so on have argued, which is the free-riding. There are countervailing arguments that suggest you could get too much, not too few, patent challenges, and we'll come back to that later.

That's an important thing I need to keep in mind.

Now, on the quantification, when we do it, simulations, what we find, again, it's still being worked on, this, but it seems to be fairly robust. Something on the order of 75 or 80 percent of applications -- of applications are made on inventions that would be developed anyway -- you know, you have the idea; the question is do you develop it -- that would be developed anyway, even if they didn't have patent protection. In other words, the patents on these are not innovation-inducing. Okay?

Out of those that apply, about 35 percent get screened out -- of the low types -- get screened out by the Patent Office. Putting those two numbers together, that implies that something like 75 percent on this argument, on these results, 75 percent roughly of patents that are granted are actually -- should not have been. That is, when I say should not have been, I mean are not innovation-inducing.

Okay, I want to just make one comment that's not in the original paper, which is patents may do other things. We know actually they do. They give access to finance. They're signals of various things. So there may be other benefits to patents, but, of course, they have to be weighed against giving patents that have dead weight -- that create dead weight loss when you shouldn't do, that is, when there doesn't increase the amount of innovation you get.

Okay. And then there will be welfare gains from several -- several different things, including frontloading fees and this new post-grant review. So let me give you just a feeling for the model very quickly. So the story is that the inventor's endowed with an idea; it could be a low type or a high type. The difference -- this is the simplified model. The low type is a patent that -- and the low type has a certain -- has a certain cost of development. And the high type has a different cost, and there's a mix in the population, okay? So lambda is the fraction of high types here.

You need to do the R&D investment to develop it. You can't patent an idea under the Bilski decision from the Supreme Court. If you don't patent it, you get some duopoly profit, and here's the one competitor, one inventor, pi, and if you get a patent, you get a premium on that. Okay.

Now, we're going to assume these two things. The first one is simply the definition of the low type. A low type is one whose development cost is below the duopoly profits even without a patent. In other words, the low type would be developed anyway. There's no additionality by giving him a patent. The high type not. The high type's development cost is above the duopoly without profit -- without a patent but below the duopoly profit with a patent. Otherwise, it's not interesting. Okay? So that's what a low type means here, okay?

Now, the patentability standard, what should the patentability standard be? The patentability standard should be -- and this is controversial, at least it doesn't seem to be appreciated by the legal scholars as far as my reading is concerned of that literature because you see things like, you know, we should give patents when -- for valuable inventions. That's wrong. From an economic point, that's just wrong. Or we shouldn't give them to low-value patents. That's also wrong.

Okay, why? Because the social planner should give a patent if it induces the innovation and the innovation would not have -- and it has social value, of course, we're assuming, that would not have otherwise been developed. Otherwise, it's just giving inframarginal rents, okay, which are costly.

And so you should only give a patent in this simple framework to a high type, okay, not to a low type. Now, you might say, well, that's fine, that's an economic definition of what we should be doing, or criterion, but is that what the courts do? No, probably not. The courts, of course, have various statutes, 35, 102, 103, 112, there are various statutes that govern the eligibility for a patent, but if you look into the literature and to the court -- to some of the court decisions, they actually say -- the Supreme Court, for example, has said in various cases what we're trying to do is this. We're trying to give patents only where they are needed to induce development. Okay?

Now, of course, they use these legal standards to do that, and so you might think that the courts make mistakes because they're using proxies. We can allow for that in the courts, but here we're assuming that they do it perfectly. But I just want to make sure that -- I want to make the argument here that this economic definition, while it appears to be at variance with what the courts are doing, it's not what -- it is what the courts say they're trying to do. Okay? In any case, it's what they should do.

Okay, fine. But, of course, the low type would also like to get a patent. Why not? It gives them a premium. So the question is how do you screen?

Now, the screening here is going to work the following way. If your type is theta, you decide whether you want to invest. To apply, you have to pay this fee, Fee A, and then you get examined by the Patent Office. Now, when you're examined by the Patent Office, if you're high type, you always pass. That is one-sided errors here. This is the baseline model. If you're invalid, you pass with a probability 1-minus-E. So E is the probability the Patent Office screens you out, if you shouldn't get it. And we call that the examination intensity. We're going to simulate the value of that. If you're granted, then you have to pay this renewal fee or this post-grant fee to activate your patent effectively, and then you move forward.

Okay, now, consider the case where there are no challenges, just to nail down the intuition very quickly. If there are no challenges, the high type invests, applies for a patent and activates -- pays the renewal fee -- if this is true, right? That's the profit minus his development cost, which he has to decide to do, minus the two fees. He knows he'll get through, so he pays both fees, if that's positive.

What about the low type? The low type always invests because even without a patent it's worth doing. And he applies if the patent premium minus the renewal fee that he'll have to pay if he gets it and activates, he goes through with 1-minus-E probability if that's bigger than the application fee. Okay?

Now, these two inequalities actually imply the following result, that means straightforward, that application fees screen better than renewal fees, post-grant fees, because the high type doesn't care because he's going to get through anyway, and the low type prefers renewal fees because he only has to pay it if he gets through. It's like you apply to Harvard to get in; if you get in, you pay the application fee, otherwise you don't. Okay, that's -- it's the same kind of argument. So the low type will be screened out if you have to pay it up-front, okay? Okay. So that's the first result.

Then what happens if -- in the licensing game? So if you get a patent, then what happens? Then there's a licensing game, and the basic structure is I offer you -- I offer a you a license contract. Let me just talk it through -- you're a licensed contract, take-it-or-leave-it offer. If you -- and I hold you down to your outside option value, which is pi, you'll get if you -- you get it anyway, and delta-

C is the decrement to profit if you don't take the license because then I'll have my lower cost from the innovation; you won't; we'll have asymmetric duopoly, okay? And so you'll suffer a decrement to your profit.

Now, if you accept, we're done. If you reject, then you can choose to challenge me or not. If you -- and that's going to be endogenous. If you challenge me, you and I each incurs a litigation cost, and in the courts in the baseline model, as I say, high types are always upheld, low types are always screened out, always invalidated, okay, in the baseline model. All this generalizes, though.

Okay. So what's -- in the presence of courts, what happens? In the presence of courts, what happens is that you get a semi-separating equilibrium, all right? You can't have -- you can't have a fully separating equilibrium, it's pretty obvious, because if you did and only the high types applied and the low types never applied, then I know that I would never challenge you because I know I'll lose because you're high type, but then a low type has an incentive to go in, so it can't be an equilibrium.

So you end up with a semi-separating equilibrium, and the one we looked at here is that the high type charges the maximum fee that it can, okay, that is, the outside option value for the -- for the competitor. The low type randomizes, here over the license fee. So with the probability Y, he charges -- he fakes it, he mimics a high type, with the probability of one minus Y, he charges the low fee.

Now, the low fee is going to be exactly the litigation cost for the competitor. In other words, I'm preempting your challenge. You know I'm low. If I charge a low type, you know that I'm low -- low type, but you don't challenge because I'm just preempting, okay, like the Barnes & Noble paid a settlement that -- that preempted them essentially, gave them no incentive to challenge the Amazon. And if you see a high type, as I say, you challenge with some endogenous probability.

Now, the one thing I want to mention is this challenge preemption is like trolling. You can think of it like trolling, allegedly, okay? Because they have a patent, it might be a low type, that's arguable, but in any event, it's low type. I extract some royalty from you, how much, just enough to keep you from challenging me. Okay?

And I mention this because the FTC, among other institutions, has been very concerned -- let me not say fixated -- concerned about trolling. I'm going to show you that that's a bad target, that you can have welfare-improving changes that increase trolling and conversely. Trolling in this model and in these simulations is endogenous. Okay. Fine, and that's what I've said here.

Now, that's the simplest model. What we've done with this revision is everything has been generalized to a much more complicated model where we allow for there to be a pair -- value -- here social value -- and cost of development, and there are distributions on both. So now it's just fully generalizing the heterogeneity in both dimensions. So you can have different -- you can have heterogenous value and heterogenous development costs that might be dependent on value, okay, because you might think that more expensive -- more valuable patents are, on average, for example, or stochastic first-order dominance might be more expensive to produce, maybe.

Fine, and everything, then, is indexed by value, so full heterogeneity. The low types, again, are just those types for whom pi, which is now a function of V, that pi is less than -- is greater than kappa, okay? So nothing changes. That's still a low type; you don't want to give a patent to him. And the high types are those where that’s not true. Okay?

Let me just skip some of this. All I want to say here before I just turn for three minutes to the simulation, which is crucial, is that you get thresholds coming out of this kind of model. And the thresholds are the following form, and that’s all I need to say, the following form. So below a certain value -- threshold value, nobody applies, fine. Then there’s another threshold V-hat where in this interval, only the high types apply, and there are no challenges, because you know you’ll lose.

Then there’s VCC for challenge credibility constraint. Now, the low types do apply. They’re above this threshold, but you will get no challenges. Why? Because not -- because you know you’ll lose, and it’s not worth -- I’m sorry, you might lose and it’s not worth -- your value is not high enough to make that worthwhile. And then above this challenge credibility constraint, low types offer -- they randomize, like I was describing, mimic or challenge preempt, and they may get challenged. Okay, that’s a characterization of the equilibrium.

Now, I want to -- I’m running out of time, but I definitely want to talk about the simulations briefly, so let me just do that. And you can talk about welfare maximization, but -- okay, so the proposition here in this case of a fully heterogenous model, you again get frontloading is optimal. I won't go into it, okay. So -- because the intuition is the same. So you still want to frontload fees.

Okay, I'm going to skip that. Okay, now, what we do next, and I'll take three minutes to do this, is we parameterize this model -- I mean, this is very stylized version of the model discussion -- in the following way. We assume and we haven't up to now, but we assume now a linear demand and Cournot behavior. Before it could be any kind of market interaction.

We use six-digit NAICS codes, so that's fairly detailed, you know, frozen peas and carrots kind of level as the market, so about 440 of them. And we extract information or construct it actually on price-cost margins, and we have the Herfindahl measure for the top 50 firms, and from this, you can actually, assuming an end firm Cournot model, you can actually infer the A and C. A is the demand parameter; and C is the marginal cost, which is assumed constant here, okay? So out of the price-cost margins for each of these markets and you can get -- and the Herfindahl measure, you can extract A and C.

Okay. Invention reduces cost by some fraction, S, and we assume that's a beta distribution, so between zero and one, and we can extract the parameters of that from using average total fact or productivity growth for each of these NAICS codes, and an R&D equation, which I'll say in just a moment -- mention in a moment -- to pin down the beta. So the details aren't important. The point is we can pin down these parameters from observed features.

Development costs are exponential with the possible dependence on S, so the magnitude of the invention might actually affect the distribution of costs of development. And then we have some other information on R&D that adjust -- this is R&D for patent applications, so we take R&D, we adjust it for patent propensity by NAICS code, and then we do various things, okay.

So and then finally, we have the litigation rate. That's the probability of being litigated -- litigated, not trial. The grant rate and -- sorry -- and the patentee win rate. Okay, we have all of this by board sector and also aggregate. And then we have a common litigation cost, which we can parameterize by -- or which we can measure from the association of intellectual property law surveys, which I'll skip that, and then we can also estimate from the simulations -- I won't go into how -- the examination cost function for the Patent Office, that is, it comes out of the simulations about what is the cost function for examining a patent as a function of the examination intensity, E. Okay, fine, so that's enough.

So the four observables that we're matching to, as it were, are grant rate, litigation rate, patentee win rate, and R&D per application. And the things that we're estimating are the examination intensity, distribution of cost parameters of development costs, and the distribution of the size of invention, okay?

And here are the results, and I'll take just one moment to -- this is the percentage -- this is the simulated -- these are simulated values for the baseline model. About 17 percent of applications are high type. These are implications of the simulation. That's shocking to me. About 35 percent of the low type applications get screened out. That means of grants, about 2 percent are low type -- are high type, okay, or 75 percent, as it were, shouldn't be granted in that sense.

Y-bar is the probability that you --

probability that you as a low type fake it as a high type. So one-minus-Y is the trolling rate, is the percentage of low types that actually preempt challenges. That's also worrisome, 91 percent.

That's the probability that you challenge, don't worry about these parameters, these are cost parameters and so on. So that's the baseline.

And, then, finally, what we do, there's some validation from various external validation or corroboration or evidence, but I'll skip that. The last thing I want to do, and I'll end in -- just very quickly, is we then do counterfactuals, right? And the counterfactuals we focus on so far -- we have others in mind -- is we frontload all fees, and we make it -- and we return the money because when you frontload fees you'll make more money subject to one condition, because all the low types are now -- are now paying it, whereas before they only paid it if they got through, and you return the money to the Government.

Then, an alternative is you frontload fees and you make it revenue-neutral. So front-load fees and then plow back in the additional revenue into the Patent Office examination. And we know how much that allows you to increase E because we estimated the cost function implicitly of E. And registration system, just said E equals zero, like copyrights, don't screen at all. And then we have various versions of this post-grant review, which basically is this internal check -- internal litigation that's much cheaper -- $350,000 roughly as compared to litigation, which is a million dollars or way more, depending on the value. Okay?

And the last slide is this one, which is what do these things do? What about these counterfactuals? Well, I want to focus -- this is the status quo, the ex ante, the baseline. I want to just focus on this. If you frontload -- frontloading doesn't -- so what does frontloading do? Not much here. But when you frontload and invest, you reinvest it, so you make it revenue-neutral. E goes from 35 percent to 45 percent. You can afford to raise it like that.

The percentage of grants goes up. Y doesn't change too much. If you have a registration -- but, so, and the welfare gains here are about 1 and a half percent from frontloading and reinvesting. And if you have a registration system, it's a disaster, welfare goes way down here, and if you have post-grant review, you get some welfare gain because you're basically lowering the cost of litigation.

So the bottom line here is that some of these reforms help; some of them don't. And we'll skip that. So the conclusion, and I'm sorry for having to rush this, but the conclusion is, first, I think we need to look at patent examination -- patent screening as beyond patent examination. It's more than that. It involves more than that one institution. And we need to have some kind of framework -- analytical framework, model if you want -- to analyze that and be able to say anything about how changing one of a combination of instruments will affect the system and screening and welfare. Finally -- and that's the main point.

And there are many other counterfactuals you could do here, interesting ones like what happens if you introduce litigation insurance, what happens if you change from the American to the English rule of legal fees, in other words, loser pays, the Actavis case about -- the recent Actavis case about pay-for-delay, that is, allowing -- restrict -- basically restricting negative fixed fees. We can do that as well. So we're going to do a number of these counterfactuals, but the main point is we need a model and we need to think about patent screening in a new way.

Thank you.

(Applause.)

DR. SCHANKERMAN: Apologize for running over.

MR. ROSENBAUM: Okay, so I think we’re going to defer questions until afterwards, but everyone is welcome to continue the conversation with Mark after the end of the conference. Is that okay?

DR. SCHANKERMAN: Yeah, sure, as you confer.

MR. ROSENBAUM: Thank you.

Okay, so with that, we’re going to move on to our final panel, which is moderated by my colleague, Miriam Larson-Koester.

PANEL: LEARNING ABOUT SUBSTITUTION AND WELFARE FROM DATA

MS. LARSON-KOESTER: Hi. So I have the pleasure of introducing this really stellar lineup of panelists today. At the FTC, we're often faced with answering a very specific question with limited data available to us. In antitrust, for example, we often have to predict how firm strategies will change following a merger, and this will depend on consumer behavior. In consumer protection, we often need to estimate the harm from firms' misrepresenting product characteristics, and so this will involve both estimating how many consumers were influenced by the misrepresentation of the characteristics and how much consumers value these characteristics.

So mapping these experiences back into the academic literature, all of these questions are fundamentally about inferring consumer preferences from data, and so we're looking forward to hearing from the panelists about how to do that best. As an FTC staffer, I hope to walk away with a better understanding of how empirical models of consumer behavior can help us get the right data and learn more from the data that we get.

So I'll introduce the panelists. We have

Steve Berry from Yale University, who you've already heard from this morning. His 1994 paper is seminal in empirical IO in mapping market shares into consumer demand, and he's continued to push the frontier of knowledge in discrete choice consumer data with work in nonparametric identification.

We have Fiona Scott Morton from the Yale School of Management. She's a former DOJ Deputy Assistant Attorney General and has work across many topics in empirical IO and antitrust.

And, finally, we have Chris Conlon from the Stern School of Business. He has worked on using experiments to estimate demand as well as developing state-of-the-art code to estimate demand.

So the structure of the panel, we have each panelist will do a short introduction to a topic, and then we will have some follow-up questions and discussion among the panel between each topic. And then at the end, we'll have time for more general questions and for some questions from the audience.

And, so, without further ado, I'll bring up Steve Berry to give the first topic introduction.

MR. BERRY: Okay, so I'm very happy to give a very short introduction here. I told my coauthor, Phil Haile, that he shouldn't worry, I was just giving five papers in five minutes, so there was no problem.

But the reason I think I can maybe do it is that I really only have one word today. It's kind of like in The Graduate when the advice is plastics, my advice is instruments. So we've known from the beginning of supply and demand estimation if we go all the way back in time that learning about demand would require instruments and, in particular, instruments for price. But in differentiated products markets, it's more complicated than that, and certainly when we wrote, say, you know, BLP in 1995, one of my very first questions from Mike Winston was, well, is this nonparametrically identified. And I just laughed and said, yeah, as if I'm ever going to prove that, but it only took 20 years. It's not that slow for me.

So if you want, we have an annual reviews piece, which has a fair amount of math in it, but it also actually really tries to get to the conclusions of what we want to do. So -- and basically the conclusion is that if you want richer substitution patterns, we start with instruments and then we want to add more instruments.

So we start with the same instruments that were used in 1929, which are instruments for price. Now, if you want super-rich price effects, you may need variation that moves the prices of different products around differentially. If you went all the way to a completely nonparametric model, you might need as many cost-shifters as you have products in your choice set if you want to -- if you want to have really completely free substitution patterns and price.

Now, what about other kinds of substitution patterns if you look at a nested logit model or you've got this other substitution parameter in the BLP model, you've got the variances of random taste? Once you think of that as the inverse demand, if we solve out for product-level unobservables, what you end up on the other side are really market shares that within group market share and the nested logit or some more complicated function of market shares in the original BLP model.

So what we really need are also instruments that move market shares, which aren't the same as the price-shifters if we want really a completely nonparametric treatment of this. So we need something like changes in the choice sets, something that moves people's choices around. One of the most natural things would be if we have access to exogenous product characteristics that move us up and down in the space of preference for different products so that we can watch where people go as the product gets better or the product gets worse. And sometimes -- we didn't call it this, but -- whoop -- sometimes people call that the BLP instruments.

So I'll just keep going. Can we have the slides back? Oh, they're over there. Okay, that's fine. I'm the only one who can't see them. That's fine.

Oh, there's one in front of me. It's the confidence monitor. I should have had confidence.

Now, you know, if you really read our completely nonparametric work, though, you might get a little -- you might get a little nervous, which is you need, like, a lot of instruments to get really rich substitution patterns. So the solutions there are just really the classic ones. Most people in practice, we don't have that much data anyway, you're probably going to put a stronger functional form on. And those functional form restrictions are going to reduce the number of instruments that you need.

Adding a cost side as in our original paper, but Chris has done nice simulations showing how important this is, adds additional restrictions, and they're more natural restrictions on the cost side because while the price of every good and potentially the characteristics of every good on the demand side, you might think on the cost side that the endogenous variable is output maybe, but it’s like my output, unless it’s a network industry or something. It’s not all the outputs. So you get many more exclusion restrictions on the cost side.

And the other thing is you might have consumer-level data. So it’s a little heroic, maybe, to get all of this out of just purely market-level data, and some microdata that matches consumer attributes to product choices are also really important.

So I think we might talk a little bit more about microdata, but I think the intuition about microdata maybe comes from the geographic example. So if you think of McFadden’s initial prediction of what BART would do where people are moving around in the space of the public transportation system or hospital demand where you get farther and closer to a hospital, so in that case, you’re learning about substitution patterns in some sense by moving people within the fixed choice set and seeing how they substitute as they move closer and farther away from different choices. And you can generalize that to other kinds of characteristics. As your family gets bigger or smaller, you're sort of moving about in the space of preferences for big cars and where do people transfer from.

So in this case, we can learn about substitution from the microdata alone, and you can do it without this exogenous variation from the BLP instruments. In the end, though, prices at the market level -- you might even define a market to be at the level at which prices vary -- and you're still going to need the instruments for price, so you're not going to get away from those initial instruments. But the microdata might get you away from these BLP instruments, which I think is potentially important.

And then I think there are all sorts of questions about how you do this once you have a functional form, and you know, how do you form optimal instruments, and how do you compute the whole thing. And, luckily, Chris has solved that all for us with this package he has up called PyBLP, which that's just my ad at the end for Chris. I'll stop there.

MS. LARSON-KOESTER: Thanks, Steve. So just as a followup question for the panel in sort of general, what can we do in terms of estimating demand if we don't have the data variation that we need?

MS. MORTON: I'm going to leave that one to you.

MR. BERRY: So taken literally, it sounds like the answer is don't, right? And I really do think that, you know, I'm sort of terrified that people say, well, you know, I did BLP, and it's like, you know, the first thing to do, it's not -- is to actually ask what's the source of variation in the data and what can we possibly hope to learn from that, right? And it's just not that different than other parts of applied microeconomics, where the first thing you should think of is what is exogenously varying and what can I possibly hope to learn from that.

And that may very well restrict the functional form that you choose. It may restrict your ambition, and at some point, you know, some things maybe shouldn't be done, but, you know, it's like any other applied micro seminar at this point, though, which is you're going to need some exogenous variation, and people are going to argue about it, and if you're an agency, you got to get something done, but you can still ask the question about, I think, what is plausible, how much variation do we have, and to sort of match what we're doing to that amount of variation.

So I don't know if you have further thoughts about other tricks we can use.

MR. CONLON: I mean, if we don't have variation in the data, we don't have -- I'll talk a little bit about what we can get from surveys and experiments later, where, like, we may not have, you know, the kind of market-level price variation that we want.

MR. BERRY: What about --

MS. MORTON: Yeah, so that's creating some data.

MR. BERRY: Right.

MR. CONLON: Yeah.

MR. BERRY: Creating more data, right.

MS. MORTON: Creating --

MR. BERRY: Right, so right. Get more data, right, is an excellent answer there.

Chris, you've done some work, I know, on, you know, sort of how to make efficient use of whatever --

MR. CONLON: Yeah.

MR. BERRY: -- variation we have.

MR. CONLON: So how do we -- yeah, how do we squeeze the most juice out of the orange is sort of a -- it's some sense, it's sort of a retro message, and so I ran a bunch of simulations on a bunch of large and small problems. And I think one of the things we found that was very helpful that I guess I didn't -- I sort of knew but didn't really know was that if you're in sort -- if you're without any cost-shifters or without -- with really weak cost-shifters is usually the bad world, right? That's the case we're most worried about. And the question is can we get reasonable-looking demand estimates from that world if all we have are access to something like the BLP instruments, like characteristics of other products and, you know, maybe cross-market variation in that.

And I think what we found was that the answer was sort of sometimes yes, and the sometimes yes was that if you had some assumption on the supply side, that is you had something that was moving costs around, even if those weren't excluded cost-shifters, those were just like characteristics in the cost function for the good, and you were willing to construct the nonlinear optimal IV, in that world, actually, we were able to get, like, pretty close to what sort of well-behaved asinthotic performance looked like. In some sense, like, we got back to the good case, even without cost-shifters.

So there's some hope without cost-shifters, but I think there’s no hope without any instruments, right? If you have the same set of products and the same characteristics and the same prices, in 100 markets, you have one observation in your data. You can’t -- sort of can’t fix that.

MR. BERRY: Right, but I think what you’re suggesting, which is always intuition, and we don’t really have a fully nonparametric proof of this, right, is if you formally add the cost side, there’s so many exclusion restrictions there that they can play somewhat the role of the cost-shifter, right?

MR. CONLON: Yeah.

MR. BERRY: Now, so, you know, and so we talked about more data, right? So, you know, really the answer is, A, you need some variation; B, you can think about the functional form; C, you can make more assumptions. Right, so this is going to mean really committing yourself to some first-order condition, right?

MR. CONLON: Yeah, you have to really believe the Bertrand-Nash first-order condition or something like that. But when you do that, you see very clearly, like you can just write it out. You get cross-equation restrictions from supply and demand on the endogenous parameters, and that’s going to be whatever the price effect is and then whatever the quantity shift -- the sigmas basically, the taste parameters, because those are the only two things that enter both equations.

MR. BERRY: It would be interesting. I don't know if I've seen this, which is how much it varies if you change the -- how much do the demand parameters change when you change the conduct assumption, right? We could try it out at some point, right?

MR. CONLON: Yeah, we could try that.

MR. BERRY: Which is, you know, some Bertrand-Nash, if you're not that well identified on the demand side, that's going to go through to the demand parameters.

MR. CONLON: Yeah.

MR. BERRY: What's the robustness with respect to different supply sides, right?

MR. CONLON: Yeah --

MR. BERRY: On the demand parameters, right? Not that you're going to get a different answer.

MR. CONLON: Yeah, that was sort of the Reynaert and Verboven exercise. What they said is they said, look, with complicated supply restrictions like Bertrand -- multiproduct Bertrand, it's really hard to back -- it’s not that hard -- to back out marginal costs from prices and vice versa, and they said if we just assume marginal cost pricing, can we construct IV that way, and I think they had, you know, some success.

MS. LARSON-KOESTER: So, yes, following up on -- mentioned by Steve is the benefits of microdata. I wonder if you could speak to how we could maybe use microdata to make up for some of the lack of variation in other ways.

MR. BERRY: Well, again, maybe I’ll start. So, yeah, I wasn’t there, but years ago, Susan Athey told me that Rita Wenbends (phonetic) had given a talk on econometric methods and had speculated at that time that market-level BLP was not identified but microdata would be. So I think that’s an extreme case -- that’s an extreme statement, and indeed we have an identification proof, so I like to say it’s not true, but I kind of know we said it, right, because it’s just market-level data, you’re trying to learn a lot of stuff from, you know, only a certain amount of data. Now, I think it really does depend on how much variation you have in the data, how much things are really getting pushed around, how many extra restrictions you’re willing to make --

MS. MORTON: How the choice set changes.

MR. BERRY: -- how the -- if the choice set changes dramatically, I think you're in pretty good shape. If not, you're not in very good shape. So, you know, what do you do when the choice set doesn't change very much? And microdata still, even though I've got papers on it, it's still slightly mysterious to me, but, you know, if you just think writing down what McAdam did, which is a likelihood with, you know, consumers in it, you can get a bunch of substitution parameters out of that. You can get nested logit parameters. You can get random coefficients out of that.

And, again, I think the intuition is that we are making some restriction. It doesn't have to be very strong, but that the consumers are moving around in some kind of space that's tied in some way to the choice set, right, that we have some variables that are sort of making one choice better for you and one choice worse for you.

MS. MORTON: Right, and we know something about those consumers, so we're able to link the consumers --

MR. BERRY: Exactly.

MS. MORTON: -- the number of children with the size of the car.

MR. BERRY: Exactly, right.

MS. MORTON: And, so, then, they're moving around in a particular way.

MR. BERRY: Right.

MS. MORTON: And buying a lot of large cars. They're never substituting to the sports car.

MR. BERRY: Right. So those kind of substitution patterns in the data, right, which are exactly -- it's exactly right -- from interactions between the people and the products, right, because, again, you can think of distance as being the easiest one, but it can be all kinds of other interactions between people and products, can show you as you change a person in a way that makes them like one product more than another product, where do they draw from, right?

What's the diversion ratio in some sense from as you move around in the space of person interacted with product characteristics, and that, I think, turns out to just be super powerful. So now we're down to just -- just needing the price instruments. And, again, you can interact that with functional form. So let's say there's just one coefficient on price in your discrete choice model.

Okay, now I need at least one good cost-shifter.

MR. CONLON: Right.

MR. BERRY: Right? That's going to move that price around, right? So you can go from needing 2J in a sort of market completely unrestricted case, 2J instruments, in other cases down to, say, one instrument in a case where you have rich microdata, you're willing to use that to trace out the full richness of the substitution patterns, and you're willing to restrict price to depend, say, on one coefficient.

MR. CONLON: Yeah, I think in practice, I think this is actually getting easier than it used to be, so, like it's not that hard now to imagine, like -- you know, one of the easiest things to do is to go -- if you're doing consumer products is to go to the Nielsen data, look at the panelist data, and just look at the correlations between income and various characteristics of products, right?

That's basically available to almost all the people in this room for some price, and so it's really easy to construct those kinds of moments from the microdata, even if our goal is to estimate on aggregate data. And so if the goal is to estimate sort of observable heterogeneity, like the demographic interactions in Nevo, like, you know, kids times mushy or something, right, that's something we could plausibly expect to see, you know, in the microdata, and that kind of variation is actually really helpful, these, like, observable interactions between, you know, price paid per surveying and income. You know, that's pretty easy to do, and that can get us a lot of the heterogeneity.

And the sort of one thing that makes that a little bit easier is that because those things are observed, you know, we can either get that across market. As income varies across market, we can get that across individuals within a market from these other sort of surveys and things like that.

MS. MORTON: Yeah. And if you have the same consumers over time, then not only do you have their demographics, you might have the choice set changing, also. And so then you really have a lot of dimensions of variation that you can exploit to identify the parameters.

MR. CONLON: Yeah, I mean, I think the real -- I mean, in some sense, if we can estimate these kinds of demographic interactions, we can almost get away without having unobservable heterogeneity, that is, you know, if income actually explains all the willingness-to-pay differences, maybe we don't need random coefficients on price that can sometimes be hard --

MR. BERRY: Right, I'll caution on that. So for years I told the story that my random coefficient on size of the car was something like the size of your family, that with a bigger family you wanted a bigger car, and then General Motors gave us this super-rich, consumer-level data, and I rushed to it to show you, you know, this strong correlation between family size and the size of the car, and it wasn't there, which was kind of upsetting.

And it turns out, of course, that we learned something else, which is that people have portfolios of cars, and a lot of people with big families buy small cars because it's a second car or they buy two small cars rather than one big car. And in that paper, we did find that income and price was very strong, but other demographic -- pure demographic interactions were not as strong as we'd hoped. I mean, so, you know, you get rural times pickup, and that's a big deal at the time, life has moved on, but at the time, greater than or equal to two kids times minivan, big effect. That was about it in terms of being able to predict things.

But what is on the other hand true and it’s not just these explicit interactions that the microdata should help you with. It should also help you get some of the -- some of the substitution in the ex space as well.

MR. CONLON: Yeah, yeah, yeah.

MR. BERRY: In other words, that you don’t have to estimate just a logit with interactions; you can estimate a nested logit or random coefficients. And those -- and that variation at the micro level helps you with that -- can help you with that as well.

But, yeah, so but panel data plays a similar role. Second choice data can play a similar role. Ranked data from a survey, if you believe it, can play a similar role as this kind of -- you know, what we call microdata, which is the one that matches the choice of the consumer to the product.

MS. LARSON-KOESTER: Do you have a recommendation for the best kind of microdata to get?

MS. MORTON: Well, it depends on your question.

MR. BERRY: Yeah, it depends on -- yeah, so -- so, I mean, okay, things that aren’t quite as good, right, but are still valuable are, you know, you have another data set that you’ve got some moments still valuable, right? But, you know, the best thing would be rich consumer interactions matched to choice sets, over time, where you see people moving within the choice set themselves, and obviously where you have a strong intuition about how these -- how these consumer-level variables are moving people within the choice set.

MS. LARSON-KOESTER: Great. So I think we're going to move on to our next introduction, which is Fiona Scott Morton.

MS. MORTON: Okay.

MS. LARSON-KOESTER: She's going to speak to learning about behavioral biases.

MS. MORTON: So I thought we were going to collude and not have slides, but I don't have slides.

MR. BERRY: The optimal response is cheating.

MS. MORTON: Yeah, I cheated, so I have no slides. I'm going to take us in a slightly different direction and talk a bit about behavioral biases and how difficult they are when you have to estimate a demand model. So search frictions have been around for a long time, decades and decades. Behavioral biases, the research on that has also been around for a long time, and in an antitrust context, that's really important to stress.

You might think why am I introducing that, you all know that, it’s because when you’re dealing with a lawyer, okay, it’s very important to say this is old, it’s known. It has a Nobel prize, okay? It’s not novel or, you know, different or unestablished or anything like that.

Okay, so they are different, however, the search frictions and the behavioral biases, because in the behavioral context, you do have these very philosophical questions about how to measure welfare, which I think introduce a little bit more trickiness. Also, I think the behavioral biases are underutilized in antitrust, and that’s something that I don’t fully understand, so I’ll talk about that a little bit.

In settings where search costs are particularly high, maybe you’ve got costly information-gathering, maybe you’ve got a very complex product, then consumer behavior is going to differ from the situation where you have low search costs. So for example, Kate Ho and I have a paper where we look at this problem in the choice of Medicare Part D plans.

So this is the pharmaceutical insurance for old people, and they have to choose an insurance plan for the next year, and they don't know what their health is going to be in the next year, they see the prices in terms of the premiums of the plan, and then there's this complex benefit schedule, which is 25 percent copay, and then there's a doughnut hole, and then there's a catastrophic region, and there might be a deductible, and each of the plans has a list of medications that's on the -- you know, first tier or the second tier with different dollar amounts attached to them. So it's quite a tricky calculation to figure out.

Now, you can go to the government's plan finder, put in your medications that you take on a regular basis and, bam, you get a sorted list from least cost to most costly. So there's a search tool that's available. Nobody uses it. And the federal government subsidizes 70 percent of the cost of this program. So the seniors are possibly cognitively challenged. They're certainly in part of the age distribution that's not as good at the web. It's a really hard problem. And they're not bearing the monetary costs of that problem.

So there's insufficient shopping. That means that there's insufficient competition on price. There's no point for one of these plans to drop a price if nobody responds to that, okay, so you have insufficient competition on price, and the benefits of privatizing this program rather than running it as just a normal government program diminish as a consequence, because why would we privatize a government program to take advantage of the benefits of competition? We don't have them because consumers aren't shopping.

So there's very little switching in the data, despite hundreds of dollars of potential savings and even more if you took the taxpayer into account. We model a rational search in that context where expected savings have to be greater than the search cost of searching to the consumer. But, of course, the search cost of searching to that consumer reflect all that consumer's life and not, perhaps, yours or my search cost of solving the problem for that consumer.

And we assume that if they search, they get the right answer. And we take this to our data and what we see in the data is that the probability of searching goes way up if you have health shock, if your existing plan has a price increase, if your existing plan has a coverage decrease. If any of those happen when -- you're less likely to roll over, so the default is you roll over and you don't shop, but if these shocks happen, then a lot more people switch.

Now, how do you accommodate this setting when you also might have persistent brand preferences? So maybe my default -- maybe I'm rolling over and that's the default, but it's also just my brand preference, how are we going to econometrically separate those. Because you have people qualifying for Medicare each year, you've got new enrollees. They don't have any switching costs. I mean, or rather the switching costs are the same for every plan. So we can infer which plans are best and what the preferences are from the entering guys and then we have the switching costs from the folks who are continuing in the program.

Okay, so we can separately identify the persistent preferences from the switching costs because of the structure of the data, and we have four years of data and the same people, okay?

And -- so what you can do is estimate this model and show that the elasticity is going to be lower than in a model without switching costs. That's kind of obvious because when a plan lowers a price, there's much less response than there would be under normal conditions.

So I'll just turn quickly now -- so that's a way you can build in a switching cost into your estimation. So that's sort of Example 1. I'll just spend a couple of minutes on Example 2. I think behavioral issues are going to be much more important going forward in terms of applications because they're going to be necessary in all of these tech -- big tech platform contexts.

Consumers don't optimize; they respond strongly to defaults. They don't search enough. So we see this, for example, if you look at the European Commission's search in Android cases, you see this showing up strongly. So the default search engine, the default browser on the handset. When something -- when a search result is presented in the shopping context, do people scroll down to the next page? No, they don't. They click on the thing that's right in front of them.

They don't invest -- consumers don't investigate a counterfactual. They don't search using another engine. They don't check if the local results would be different if they used a different shopping service, so they don't know the quality penalty they're paying from lack of search, and that then enables that to be an equilibrium behavior, okay? And they might need a shock like the kind that we used in the Medicare context, like a health shock, to cause them to find out.

We see that platforms take advantage of framing effects. The space taken by the box at the top, all of this dark patterns literature, there's a really nice Stigler report on dark patterns, how you can get people to buy an insurance product they don't really want by making it really hard for them to escape buying it. And so the consumer lack of search makes the ordering really valuable; therefore, the search engine is going to sell the ordering, and, therefore, providers are going to buy the ordering. So that's a totally rational business model in a world where framing and consumer behavioral biases lead to this kind of behavior.

So you probably want to model those choices as a function of the results ordering, not of some utility maximization across all the choices in the choice set, okay, because that's not really what's happening. So I will stop there because I'm out of time, but I think that that's just an extremely important reality to be modeling when we go to these settings where choice sets are not really the same as, say, a grocery store where everything's on the shelf and maybe one's a little lower and one's a little higher, but they're all visible right there.

Okay, thanks.

MS. LARSON-KOESTER: Thank you. So you mentioned sort of the nonsearch costs affecting how competition plays out in a market, and I'm just wondering if the panel can speak to sort of what circumstances do we know -- or how can we find out if behavioral factors are something that will be important to consider.

MS. MORTON: Yeah, I mean, I think -- I don't know if there's one single test that says, okay, here's a behavioral factor. I think it's the economist knowledge of the choice environment, of the search environment. Is it the case that there's a tool that everybody's using that's ranking something at the top, that's the case with a lot of digital applications.

In the case of Medicare Part D, the old people are not using the web, and so there isn't a tool, and what's -- what does search look like in that environment? I think we have to know the institutional details of our market, and then we have to be attentive to the literature. I mean, you can't -- you can't read something that says competition is a click away and reconcile it with $12 billion to be the default search engine, right? If the default search engine wasn't doing something really valuable, nobody would pay for it in that way.

So I think the model has to reconcile all the different facts that the economist knows about that marketplace.

MR. BERRY: I mean, I don't know how much you've thought about this, but one of the scary things for me about the whole framing literature is that, you know, if framing matters, it means that relatively small changes in a choice set or something could have relatively big differences in the outcome, right? And it almost seems to lead to kind of -- almost to kind of discontinuity perhaps, you know, maybe not hundreds of dollars' worth, but that relatively subtle things are going to make a big difference.

And, I mean, is that a big challenge here that, you know, there's sort of no -- not like we're looking at the price or something like that, you know, sort of what --

MS. MORTON: Right, it's not as smooth, yeah.

MR. BERRY: -- what the frame is, right?

MS. MORTON: Yeah, yep.

MR. BERRY: Yeah, I'm a little worried about sort of a lack of smoothness and a lack of -- you know, it's always hard to go to the next example from one example that that makes this much more difficult.

MS. MORTON: Yep, yep.

MR. BERRY: So I find it a little frightening.

MS. MORTON: Well, I mean, it's true, it's not going to be smooth, but I agree with you that if you have a setting where the consumer can see the first three ranked options on her mobile device and nothing else, a pretty good model might be those three. And if something changes in the algorithm that puts a different three in, then maybe the choice is now the consumer's choosing among those -- that different three.

I think that's -- I mean, I would start there. I'm not sure where it would lead.

MR. CONLON: Yeah, I mean, I think having -- in general, I think having the data on the observed search process or the institutions that govern the search process is really important.

MS. MORTON: You walk up to a vending machine, and there are the things --

MR. CONLON: I know, I liked vending

machines for so long because you know exactly what's in the choice set, and that's really well observed to consumers, but other people like Ali Hortacsu and coauthors have looked at car insurance, where they have data on here are the ones you saw, here are the ones you got quotes from and so on. And those -- you know, in that case, I think it's possible to estimate, you know, what my marketing colleagues would call the search funnel of, like, the things you're aware of, they things you're considering, and then the things that you choose.

I think the test that I think is, like, I find hopelessly hard that people sometimes try to do is to estimate sort of unobserved consideration models, where we see all the products. We don't know which ones are considered, and we don't have any data on that, and then we try to figure out what the consideration set is, this latent consideration set.

And I think usually what it's standing in for is just that some products are more similar to others and we can't really tell consideration from preference in a lot of those worlds, you know, unless -- but the welfare implications are different, right? If I could just tell you about a product, now if it's really you're not considering it and you would like it, then there's going to be a positive welfare game from just, you know, informational interventions.

MS. MORTON: Yeah.

MR. BERRY: So you made a connection that I thought was unexpected to me, not to you, which was, you know, we were talking about the benefit of the supply side in demand estimation, and you suggested the benefit of in some sense the supply, in other words, the supply of, you know, the rank -- the auction or whatever that gives you the rank. Is there work that actually really incorporates the price paid by the firm, the value paid by the firm? I mean, I know you came up with some examples for us, but --

MS. MORTON: To be at the top of the list.

MR. BERRY: Yeah, that we sort of -- rather than trying to infer it from consumer behavior, we actually infer it from the behavior of the firm. In other words, the firm is telling us what matters.

MS. MORTON: So I do not know of such a paper, but that would be a great paper for somebody to write. Now, you'd need to know how much the search engine or the bottleneck that was doing the framing for the consumers was charging. You need to know those prices, so winning bids or contract prices or something, so that is -- I don't know of data like that.

MR. BERRY: To me, it seems like a lot of those papers are focused on just the revenue to the platform or something like that, whereas you're suggesting, I think, something much more interesting, which is the actual, you know, value of the frame itself.

MS. MORTON: Yeah, those things should be related.

MR. BERRY: I agree, right.

MS. MORTON: Yeah.

MR. BERRY: I'm just saying that's not -- that's often not presented as that being the research question.

MS. MORTON: Yes, correct.

MR. CONLON: Yeah, I think getting data from the ad exchanges is going to be the hurdle, right? It's like --

MS. MORTON: We need you to do that.

MR. CONLON: -- yeah. Yeah, thanks. They're, like, super secretive, and then if you got the data, it would be, like, probably more data than we could store on a computer.

MS. LARSON-KOESTER: Well, I think we should move on to the next introduction, which is Chris

Conlon.

MR. CONLON: Do I get slides? Oh, great.

Okay, so what I'm going to ask is I'm going to ask, like, what is it -- or I was told to answer the question, what can we learn from experiments. And, so, a lot of this is going to be informed from some work with Julie Mortimer, most particularly this paper that isn't obviously about experiments from the title. The paper is called "Empirical Properties of Diversion Ratios."

So the first thing I think you need to ask when you run any kind of experiment, you know, in an antitrust context, you're going to use data from an experiment, is I think you want to ask what is the object you're actually trying to measure and, you know, what does that look like. And, so, what I've done is I started here just by saying, like, look, there's a lot of different ways we could measure substitution, and I've sort of focused mostly on different measures of diversion ratios here.

So the first one is sort of the classic one that you would get in Farrel and Shapiro. It says, I raised the price of Good J; some people leave Good J; and I want to measure the fraction of people who substitute to Good K. All right, so that's often in something like a UPP calculation. I think this is essentially what people have in mind when they're talking about measuring substitution.

We could also think about a different context. We could think about instead of perturbing the price of the first good, you could imagine instead we could perturb the quality of the first good. And there might be markets where that's going to be the available variation, or maybe that's closer to the experiment we could run, you know, we could see somebody makes the size of a bottle of ketchup smaller or something like that, and the quality is going down, and we could see how that leads to -- traces out substitution.

The third one, the thing that I've labeled ATE there, what that is is that's just saying, like, suppose I took a product completely away from consumers and I removed it from the choice set, right? So you could imagine, these are experiments you could run, and these are the kinds of experiments we actually ran in vending machines. We actually tried running price experiments first, and we mostly failed because it was -- you know, nobody responded to five-cent price changes in a way that we were able to measure effectively at the frequency we had in our data, but, you know, if we took away the best-selling products, then it was actually something you could actually maybe hope to measure.

The final thing I put up there for fun is, like, the logit. And I put up the logit because if you sort of have just diversion proportional to share, it turns out all three of those measures that I wrote are all going to be identical in that world, but remember, you're predicting substitution with not a no-parametric -- not a nonparametric model, but rather a no-parameter model, right? And sometimes you're not estimating anything.

And, so, the other thing, you know, experiments can tell us about is they can tell us about welfare, right? And so what I did is I just put up, like, the logit sort of a random coefficients logit version of consumer surplus, and it turns out that, you know, what you get is you get, like, as I change prices, what matters for consumer surplus, at least sort of the best approximation, is how much the outside good share responds, right? So how many people are switching from buying any of the products to buying the outside choice, right? And that's going to be true if we change prices or if we change quality and also if we change variety, right?

And, so, you know, these sorts of calculations, actually they're not -- the math is really easy in a logit. It turns out that, you know, these calculations are more general, like this is what people in public finance do all the time. They say, I tax Good 1; I see how much -- I tax Good 1, maybe that's alcohol or cigarettes, and I look at how demand for the entire category responds. It turns out that's a pretty close first-order approximation to welfare for a broad class of models, right?

The other thing -- can I go back?

MS. LARSON-KOESTER: Use the red button.

MR. CONLON: Use the red button, okay.

The other thing I guess I should point out is that -- well, there's two things. One is that we don't always observe the outside good share, so that's something that's often coming off of an assumption. So it makes welfare a little bit tricky, and I'll talk a little bit about how we can resolve some of that. But I think that's good. I'll move on from there.

All right, so what we can do, then, is we can actually sort of, like, try to plot the objects that I talked about that one of my plots did not make it. We can plot the objects that I talked about, so that blue line is, like, as I trace out these, like, small price increases and I continue to increase the price of Good 1, I can measure substitution to Good 2. What the red line denotes is the same thing, but where I trace out -- as I change the quality of Good 1; and the dotted line there is, like, if I took Good 1 away completely how would people substitute to Good 2.

And, so, I've sort of just marked off like a 5 and a 10 percent price increase, and the X axis is like the fraction of sales of the initial product that are still remaining as I raise its price or reduce its quality. And, so, what's going to happen is whenever I sort of manipulate the price or change the quality or remove the product completely, I'm going to basically be tracing out a different line, and I have to make sure -- you know, this is sort of similar to what -- you know, what the program evaluation folks told us, that, like, different instruments identify different effects. And we have to be a little bit careful to make sure, like, we're getting the effect that we want.

And so here's the kinds of experiments that I think, like, people at the agencies -- both here and elsewhere -- would do. One is, like, you know, what happens, what kind of experiment, and maybe we see that a firm in its course of business tried out a small price change. You know, one of the challenges that, you know, a lot of times it's hard to measure anything for a very, very small price change, often because our data are noisy, that just demand is moving around.

The other thing that they do, and I mostly associate this with the U.K., which is why I said where would you shop if we closed this Tesco, because they love to run consumer surveys where they stand people in front of a Tesco and say, where would you shop if we closed this place. And it's clear what that's not providing information about is, like, small price changes. That's providing information about what would happen if we removed the product from the choice set, right?

And then, you know, the stuff that I've worked on, you know, obviously would be -- it would have been much easier if we did it online, where what we did is the exercise Fiona described, which is we sort randomized search results to consumers on Amazon or eBay or something, but we were dumb and we decided to do this in practice with actual vending machines, where we had to pay people to take away candy bars and hide them and things. And so -- but you can do sort of those kind of product removals or stuff like that, right? And you could think about short-run, stock-out events as sort of representing a quasi-experiment, that sort of once we condition on some things, it's going to behave as if it were random variation.

The hard part is, I think, like we need to know what's the object we wanted to estimate in the first case, and oftentimes the experiment gives us one of the other objects, right? We have this great experiment on second-choice data, but I want to know what happens when I increase my price by a small amount, right? Or I see, you know, maybe I do see a price change or, you know, some weird thing or something gets hit with the tax, but what I really want to know about, what would matter for the market, is what happens if actually we closed this store down, if we did remove the Tesco, not if we, you know, raised sales taxes 5 percent or something, right?

And, so, usually, I think about UPP calculations, at least if you sort of believe the derivation in Farrel and Shapiro as being about starting at the small price change world about the current market price, whereas if we think about something like hospital -- you know, hospital or insurer cases, oftentimes they're going to be interested in this willingness to pay object, which looks more like the consumer surplus or welfare calculation I showed you. And in a sense, that’s really about second-choice data or variation in the assortment.

I think the unfortunate thing is it’s sometimes easier to learn about the first case by product removals and the second case we don’t -- you know, sometimes we see hospitals close or insurers exit the market, but oftentimes we’re trying to learn about those from small price variation. So it’s a little tricky, right?

And, so, just my last slide here, you know, can we do antitrust with experiments only and without empirical models? You know, yeah, I sort of would love to live in this hypothetical world where what would happen would be, you know, the merging parties would come to the agency and they would collectively design an experiment that would be run by one of these consulting firms, but I think that probably is not going to happen anytime in my lifetime, and so, you know, what are we left with?

I think if you read sort of the guidelines in 2010 and sort of the literature around it, I think Farrel and Shapiro were sort of hoping and we could sort of see diversion in normal course of business, you know, that this would just be like a number in an email or a spreadsheet or something like win/loss data or, you know, cell phone porting stuff. And there's lots of cases like that, and I think, you know, there may be cases where that's possible.

I'm a little skeptical we're always going to see the object that we need, and so I think often what we're going to be stuck with is we're going to be stuck with trying to use our experiments in addition to our models as sort of, again, extra moments or extra information that we may want to match.

I think there are still some -- a lot of open questions about how do we combine these things and how do we balance experiments and observational data. You know, if I have 100 million observations from my observational data and one week of experiments, you know, there's a sense in which my model may not really care very much about that one observation of experiment. I think we need to think about how we want to balance that stuff. So that's --

MS. MORTON: Do your ad. Don't you have an ad slide?

MR. CONLON: Oh, I have an ad slide. Yeah, I was going to save that for the --

MS. MORTON: Oh, oh.

MR. CONLON: -- yeah. I'll say that later.

MS. LARSON-KOESTER: Thanks, Chris.

So following up, I know you talked a little bit about sort of what object are we actually measuring with experiments, but I'm wondering if the panel has thoughts on how we should assess the external validity of an experiment.

MR. BERRY: Sounds like a no, but I think what's useful about what Chris said is, of course, that he wants us to focus first on what question we're answering, which has to be part of the way there. And, of course, there's a very strong connection between the different sources of experimental variation and what they reveal in our early discussion of instruments and what, you know, price instruments versus, you know, sort of substitution pattern instruments. There's a very strong connection there. So to be careful about the -- to be careful about the question you're answering, but as far as external validity, I think it's really hard, actually. I think that the development RCT literature has this problem. People scale up and weird things happen. So I think we should be pretty cautious about it.

I worry, for example, about, you know, elicited preferences, for example. So standing outside the Tesco and saying where you would go -- okay, I think that's not so bad, right? How much would you buy if the price were 10 percent higher, I don't believe at all, right? And then the question is --

MR. CONLON: I mean, I think that's why the Competition Commission stopped asking that question in the U.K.

MR. BERRY: Right, and years and years ago, I was actually working on an antitrust case for something else, and they actually ran people through an experimental supermarket, having raised the price of one good by 10 percent, right? And they ran many people through the supermarket, and they were going to get the price elasticity out. They were very happy with themselves. And, you know, people didn't change their behavior at all.

And you could say, well, okay, it's -- you know, price is perfectly -- you know, demand is perfectly inelastic, but I don't believe that. So, I mean, I think the other problem with these experiments, you have to come back to the framing question. People think they're in an experiment.

MS. MORTON: Yep, yep. I would also say, I mean, external validity of an experiment in one place to something else is, I think, very counternormative to what we do in IO, where we think that the setting, the kind of people, the kind of consumers, the kind of product, the product, you know, production function, costs, informational environment, is really quite specific, and you could get a really different answer if you changed one of those things, so certainly I think external validity to other stuff should be treated very cautiously.

MR. CONLON: Yeah, I mean, I think we spend a lot of time, right, like what is the relevant market and, you know, where is this effect going to matter. And I think -- I mean, that’s sort of our version of external validity here, right, understanding how to extrapolate from what data we have and what model we have to like in this particular part of Texas in this market that this is where we’re worried about the largest price increase or something.

MS. LARSON-KOESTER: So also following up on something Chris mentioned, I wonder if the panel has thoughts on sort of best practices for incorporating other data sources like costs or margins or survey data into a demand estimation.

MS. MORTON: And you make more moments if you can.

MR. BERRY: Yeah, I mean, and the margins are, you know, in some sense an even better version of the first-order conditions, right, if you believe them, if you believe they're marginal cost.

MR. CONLON: Yeah, I mean, I think the challenge is we don't always know -- you know, accounting data may not give, you know, economic partial cost -- that's usually the big caution.

MS. MORTON: That's actually a big difference between academics and enforcement. When I was doing this, there was a lot more use of accounting data than academics would ever allow their graduate students to do. Is that fair? Yes.

MR. BERRY: Yeah, no, but you can see why, right? Because that's actually extremely powerful information, and so, you know, the approximation there in a short project may be worth it, given just how much information is in a market.

MS. MORTON: Well, and it's also subpoenaed accounting data, which might be a little better, maybe, than public accounting data. I don't know.

MR. BERRY: Yeah. In general, though, I think one advantage of writing down a relatively complete model of your situation is that it then tells you how to incorporate all the other information that you might get, right? So, you know, if you have a moment -- if you have a margin, right, it tells you what to do with it, and I think that's just really useful.

MS. LARSON-KOESTER: So I have just a few more general questions before we move to audience questions. Does the panel want to talk a little bit about best practices in general? So what are some key choices?

MR. BERRY: Chris does.

MR. CONLON: Yeah, can you put up my slides?

MR. BERRY: Chris does.

MR. CONLON: Can you put up my -- yeah.

So, yeah, I mean, I think, like, you know, what are the best practices. So what we tried to do -- I'll show you the ad here -- is we tried to sort of do them all, and so here's, I think, like, what I would tell a student to do or what I would try to do myself. I think, like, what are the objects we're going to need in a model. I think the most important objects are going to be we want some heterogeneity in the taste for a constant or an outside option because, remember, that's what's going to drive our welfare from that expression I put up before. And often, you know, the outside option is a thing we -- the size of the outside option is the thing we have the least data on to start with, so we want the most flexibility in that substitution so that we can at least -- even if we're missing the level, we can get the substitution right. That's going to give us welfare. And then similar for price, obviously we want as much -- you know, we want heterogeneity and sort of willingness to pay sort of the next thing. So that's sort of our objective of what a model should have at the bare minimum. Otherwise, we're basically just doing everything proportional to market share. We're not using any data at all.

So the next thing is, like, we should have instruments for both the prices and the random taste, as Steve talked about this. What would I do today? J.F. is here, so I would say I would follow his recommendation for generating sort of BLP-style instruments, you know, how to use characteristics of other goods in the right way, and then once I did that and estimated demand, I would probably construct the approximate optimal IV, sort of in this Chamberlain or sort of the Reynaert and Verboven sense.

You know, what I would do is if I believed I had supply conditions, I would impose them. That is, if I knew static Bertrand-Nash was what I believed firms to be doing, I would do that. If I knew they were all colluding, I would sort of impose that. And then if I could sort of collect extra micro-moments, like from, you know, survey data or other data, I would do that.

And so the shameless plug is, of course, like all those steps are hard except that I just spent a year with a student making them as easy as possible, so here's my shameless plug. We have this software, pyBLP, that will let you do all that stuff sort of, you know, in a single line, you know, one at a time. And so if you have Python installed, you can just do "pip install pyblp," and you will be able to do this sort of -- you will be up and running in, like, less than a minute, and then documentation is, like, super long and super detailed. And when it doesn't work, you can email me, but sort of here's my --

MS. MORTON: You are going to regret saying that.

MR. CONLON: -- real ad --

MR. BERRY: Don't say that.

MR. CONLON: -- which is this is the original BLP paper, and we did it in nine lines, it looks like, and so that's the full model in BLP with linear demand and nonlinear random coefficients and supply and demographics and all that, and you can see that's what it looks like when you just sort of load it, and then if you want to just estimate, you just type "dot solve," and once you've done that, then you can compute elasticities and diversion ratios and consumer surpluses and evaluate a merger with different ownership and then compute the optimal IV and resolve and everything. And, again, you know, nothing is more than a line. And, so, the hope is we can get people to, you know, use at least one or two random coefficients and we can move hopefully -- my dream is to move us away from the logit world, right?

MR. BERRY: Okay, but let me say it's like late-night television, but there's more. They have basically, I think, all of the published and folk wisdom here about how to compute different things, kind of, you know, both in the accompanying paper, you know, how do you solve this, how do you solve that, how do you deal with the exponent -- I mean, a hundred different things in here that they've just really put in one place. So, you know, it's like -- I haven't used it yet, but I've talked to a bunch of people who have, and my impression is it slices and dices it and makes you toast.

MS. LARSON-KOESTER: So pushing us a bit

beyond, you know, the known world, does the panel want to talk a bit about how we could use machine learning to improve on demand and welfare estimation?

MS. MORTON: I definitely do not.

MR. CONLON: Sure, let me start with that. So I think there's some machine learning things we can do today that you could do -- you know, you could do right away, without learning any machine learning or without really changing what we do. And I think, like, one of the things they do a good job doing is separating between the parameters we actually care about and want to interpret and the parameters that are basically just there to improve fit, these sort of what we would call nuisance parameters in econometrics.

And what I'll -- the bold claim I'm going to make is that almost all of the sort of random taste and substitution parameters in a BLP model are nuisance parameters, that is, nobody cares what the coefficient on horsepower actually is or whether there's, like, a variety of tastes for mushiness of cereal, right, we just want to get the right diversion ratios and the right substitution patterns out.

And so I think one of the things that I've found was helpful was, like, you know, we did this stuff with cereal where we downloaded 40 pieces of nutritional information and lots of product characteristics and advertising data and all kinds of stuff for cereal. What we tried to do is we basically said, actually, what we're going to do is we're going to project it down into three principal components that are going to explain 90 percent of all the characteristics that make cereal different. And it's much easier to estimate, you know, random coefficients on three principal components than it is on 37 pieces of almost perfectly collinear nutritional data, right? So that's one thing we could do, you know, today without, you know, doing much.

I mean, the other thing is we could do similar things with -- you know, using either principal components or LASSO regularization or something on the set of instruments that we put in, right? And, so, lots of people in econometrics have discovered maybe I don't need a thousand instruments; maybe I can select a hundred that are actually, you know, strong or that explain all the variation in the thousand.

MR. BERRY: Yeah, so I agree with all that, but let me give the counter case of things that people are doing that I think are right. And they mostly involve not making the distinction that Chris just suggested, which is you would not want to have a big machine learning thing predict demand and let it throw price out because it was, you know, predicting price from three other characteristics, right? So you really have to be careful about the things that you insist are in versus maybe controls or something or pure dimensionality reduction thing that you want it to do.

And then the other thing -- another similar thing is -- I mean, actually, there's a paper out there right now which basically says the problem with demand estimation is that, well, there's nonobservables, so we're just going to use machine learning to predict things and then we don't have to worry about the unobservable. But, I mean, that's obviously wrong because you need to preserve the variation and supply. You need to preserve the excluded variation and supply. So you still have to think about the model. You still have to think about things that you definitely want to exclude, like pure supply shifters, and you still have to think about things that you definitely want to include, like, say, price because you wanted the price elasticity.

I think, though, for sets of controls for the functional form of instruments, for reducing a high dimensional space in the first place, for turning text maybe into characteristics and variables, I think there's a lot of fun stuff and the correct stuff people can do.

MR. CONLON: Yeah, I think -- I mean, I think the stuff that's less available today that's probably worth thinking about is thinking about, you know, one of the takeaways from the machine learning literature is, like, you should, you know, estimate your data many times -- you should estimate your model many times, and often you want to do things like reweight the observations you can't explain or something, like put more emphasis on fitting the things that are really hard to fit.

And, so, some things like that and some things like if I -- could I take the prediction from two models and average them, I think those are the cases where if I had to forecast what we'll see in the next few years, people trying, I think it will be stuff like that.

MS. LARSON-KOESTER: So we have a lot of fun things to discuss, but I want to allow some time for audience questions, so if anyone wants to ask a question.

MR. GANAPATI: So often we might have only data in a subset of choices, so think of a case where at least at the first request we have data from the two merging parties but not a bunch of the other competitors. What would you do in that case?

MR. BERRY: Okay, it would be better if you have some data on everything else, right, because then you can get into these situations where you're trying to fit, you know, different levels of detail of data. For example, if you had market data for everyone but only the microdata for the subproducts, then I think you're actually in pretty good shape. Otherwise, I think it's really hard because, you know, you're going to end up putting the rest of the world into the outside good or something, and you're probably really going to miss some important substitution, I would think.

MR. CONLON: Yeah, I mean, my experience with the paper with Julie is we've tried sort of constraining ourselves to this exercise. Could we estimate diversion using only data from merging parties, and it was like a colossal failure for us. You know, it turned out what we really needed -- or the assumption that turned out to be really valuable was that the diversion ratio summed to one, that you had to sort of see the whole market and see where everybody was going.

And I think, yeah, there was this hope in the 2010 guidelines that maybe if we moved to diversion we could only look at data from merging parties, and it's probably possible in some cases, but, you know, the two that I've tried, it didn't work.

AUDIENCE MEMBER: I was wondering if you could talk a little bit more about welfare in the behavior bias context in the sense that if I go to the grocery store and I'm -- there's cereal on different shelves, and basically I'm just looking at the top shelf that's right in front of me, I'm going to get a very -- if I just estimate standard demand, I'm going to say, well, people must really like all the cereals that are on that top shelf, but really they're just buying them because they're on the top shelf.

So how should we -- yeah, if you could touch on how to think about welfare in that context and especially how agencies are to deal with that when we're thinking about particular applications.

MS. MORTON: Yeah, so welfare, when you've got these behavioral biases is a deep philosophical problem. Am I talking about the consumer's future self or herself at that moment? If we know that we're going to -- if I know I'm going to get addicted to cigarettes and consume more cigarettes, are you measuring the welfare of me consuming more cigarettes or my long-run self?

That's really hard. I think one of the things about the cereal aisle is that intuitively the cost of reaching down a foot to pick a box of cereal from a foot-lower shelf than the one that's at eye level is we don't see as strong and effective framing as we do when we go to a mobile -- device like this or something even smaller where there's an ad taking up two of the four square inches and then -- and then some result below that. So the strength of the framing effects is something that I think is really important to try to measure because in the cereal context, we might want to have the whole shelf be the choice set, and in the online context, we might not want everything that's ten pages down to be in the choice set. And that's not something I think we can intuit a priori.

I think we really need to have some evidence about how people behave because it's so sensitive to the way consumers -- to the way consumers operate. And you can also imagine in the digital context that

that framing would adjust to what the platform is measuring your blood pressure to be on the Fitbit and whether you're in the middle of your commute and you normally get home by 6:00 and whether -- you know, all the other information that the platform knows about your -- that might be an input into your bias at that moment.

So you've got the ability of the platform to respond in real time to what it thinks the behavioral biases it's facing are, and the supermarket has to pick some display for the shelf that is kind of some, on average, good thing that will work for most consumers all day. So it's really -- you would expect the platform to do a better job at extracting surplus in the supermarket.

MS. LARSON-KOESTER: I think we are about out of time.

MS. MORTON: Okay. Thank you.

MS. LARSON-KOESTER: Thank you to this fantastic panel.

(Applause.)

MR. ROSENBAUM: So thank you very much to our panelists and our moderator. Thank you all for joining us at the conference, and the conference is now over, but we hope to see you again next year.

Thank you.

(Applause.)

(Conference adjourned at 12:44 p.m.)

CERTIFICATE OF REPORTER

I, Linda Metcalf , do hereby certify that the foregoing proceedings were recorded by me via digital recording and reduced to typewriting under my supervision; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were transcribed; and further, that I am not a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of the action.

Linda D Metcalf

LINDA METCALF, CER